


ICASE

SEMIANNUAL REPORT



April 1, 1996 through September 30, 1996

*NASA Contract No. NAS1-19480
November 1996*

*Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA 23681-0001*

Operated by Universities Space Research Association



*National Aeronautics and
Space Administration*

*Langley Research Center
Hampton, Virginia 23681-0001*

DTIC QUALITY INSPECTED 4

19970203 059

CONTENTS

	Page
Introduction	ii
Research in Progress	
Applied and Numerical Mathematics	1
Fluid Mechanics	21
Applied Computer Science	38
Reports and Abstracts	58
ICASE Colloquia	61
ICASE Summer Activities	65
Other Activities	72
ICASE Staff	73

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U. S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, fluid mechanics, and computer science in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Theoretical and computational research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, combustion, and acoustics;
- Applied computer science: system software, systems engineering, and parallel algorithms.

ICASE reports are considered to be primarily preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period April 1, 1996 through September 30, 1996 is given in the Reports and Abstracts section which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-19480. Financial support was provided by NASA Contract Nos. NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

SAUL ABARBANEL

Absorbing layer boundary conditions in Computational Electro-Magnetics (CEM)

Computing the electro-magnetic field created by a radiating body is a challenge to time dependent finite-difference methods. Because usually the typical wave length in the field is much smaller than typical body dimensions, a larger domain requires an inordinately large number of computational nodes. Shrinking the domain raises the question of how to impose numerical (artificial) boundary conditions so that outgoing wave will not be reflected into the computational domain and contaminate it. Recently (1993) Beranger has proposed a "perfectly matched layer" approach. The layer is adjacent to the artificial boundary and is supposed to have such properties that plane-wave will decay irrespective of their direction of propagation. The field inside the layer need not be described by the free-space Maxwell equations. The overall solution must, however, be continuous at the boundary between the internal computational domain and the absorbing layer. In order to fulfill these requirements, Beranger used inside the layer a set of Maxwell like equation, but with the magnetic field eq. being split. More recently Ziliokowsky (1996, unpublished yet) has constructed, on a physical basis, a set of equations for the absorbing layer which includes an equation for a "current". This current also serves as a forcing term in one of the electric-field equations.

We have conducted a mathematical analysis of both methods. We found out, to our surprise, that the Beranger technique is only weakly well posed, and the solution-error for the semi-discrete problem, as well as for the Yee finite difference algorithm, grows linearly in time. We were also able to actually solve the set of equations proposed by Ziliokowsky. While his formulation is strongly well posed, the solution in the absorbing layer does not decay for all propagation directions. The cutoff angle for non-decay depends on the resistivity postulated in the layer and the frequency of radiation. For very high frequency only grazing wave will not decay. We were able to show that Ziliokowsky's equations, constructed from physical considerations, can be viewed as a special case of a larger family of equations. However, that larger family of equations may also include more equations, thus raising the cost of computation. In at least one such variant, the cut-off angle can be made much smaller at high and moderate frequencies, thus decreasing the overall reflection coefficient.

Future plans include a search for the "holy grail" of perfect matching with strong absorption at all angles of incidence. Generalization to 3 dimensions is contemplated. Numerical experiments, being conducted by a graduate student, will continue in order to verify the analysis.

This work was done together with David Gottlieb of Brown University and would have been impossible without the use of the large blackboard in the ICASE Conference Room.

BRIAN G. ALLAN

Fluids, controls, and rigid body dynamic coupling

In nonlinear flight regimes, the interaction between fluids, body dynamics, and controls can critically affect the performance of an aircraft. By coupling these disciplines one can computationally investigate the problematic nonlinear portions of the flight envelope. The application which will provide the impetus for this work is the controlled landing of a high-performance powered-lift aircraft. This flight regime is currently problematic due to difficulties in obtaining an adequate representation of the stability derivatives. In turn, use of a poor plant description to design the control system can lead to the loss of aircraft and pilot. To capture the nonlinear nature of the flow field, higher-order models need to be considered. In this investigation, the Reynolds-averaged Navier Stokes equations are coupled with Euler's equations of motion along with an active control system. The results of this research will begin to provide an aircraft designer with a means of computationally prototyping the control system in problematic nonlinear flight regimes avoiding costly errors.

Towards the goal of a simulated controlled landing of a powered-lift aircraft, the control of an airfoil was investigated. Numerical simulation of a two-dimensional airfoil controlled by an applied moment couple in a linear flight regime was initially studied. This simulation coupled the Reynolds-averaged Navier Stokes equations with Euler's equations of rigid body motion with an active state feedback control system. A nonlinear system was then studied by adding jets near the trailing edge of the airfoil. Robust control design techniques were used to develop a controller for tracking altitude commands. The controller was designed using a linear state space model developed from input/output data generated by the nonlinear coupled system. Simulation of the controller with the coupled system showed good altitude tracking performance. However, the controlled system produced mild oscillations in the attitude, pitch rate, and vertical velocity of the airfoil. These two-dimensional simulations demonstrate the methodology of evaluating an active control system for a rigid body aircraft in nonlinear flight regimes.

This work was done in collaboration with Professors Maurice Holt and Andrew Packard of U.C. Berkeley. They will extend this study into three dimensions for the controlled simulation of an aircraft during landing.

EYAL ARIAN

Airfoil shape design

Our goal is to develop efficient, simple to use, algorithms for shape optimization of airfoils.

Our approach is based on the following principles: (1) Use pointwise (infinite dimensional) representation of the airfoil; (2) Derive the adjoint equations in the continuous level and then discretize (to compute the sensitivity gradients); (3) Precondition the gradient direction by approximating the inverse of the Hessian in the continuum level; and (4) Estimate the step size by the above approximation thus avoid a line search.

We have implemented successfully this new approach on optimal shape design of airfoils under subsonic, transonic, and supersonic flow speeds using Euler equations (with TLNS3D). We find the new method computationally efficient and simple to apply. We intend to further develop this method.

This work is done in collaboration with V. Vatsa of NASA Langley Research Center.

Coupling issues in MDO

The objective of this work is to develop quantitative tools to estimate how tightly coupled is a given system of coupled PDEs with respect to analysis, sensitivity, and optimization computations. This work is relevant to algorithm development for problems in multidisciplinary analysis and optimization.

The approach is to analyze the coupling with local mode analysis in the PDE level or with linear algebra in the discrete level. The predictions of the analysis are then verified numerically on simple model problems. Preliminary numerical results were obtained with G. Hou of ODU for static potential flow over a plate.

Future plans are to further investigate the applications of the above methodology to MDO problems.

Local mode analysis of the Hessian for optimization governed by PDE

This work is concerned with the analysis of the eigenvalue distribution for Hessians resulting from optimization problems governed by PDE. The goal is to construct efficient preconditioners for such problems.

Local mode analysis of the Hessian has been done for optimal shape problems governed by inviscid flow and for steady state aeroelastic optimization problem. We are working on the extension of the analysis to viscous flow.

This work is done in collaboration with S. Ta'asan of Carnegie Mellon University.

MICHAEL BOOTY

Microwave induced ignition of a combustible solid

Microwave heating is being used in a number of material processing applications. One example is the reaction-initiation and control of a chemically active combustible solid mixture, which has applications to the process of self-propagating high-temperature synthesis (SHS). The formation of a desired reacted solid product can depend on details of the inert warm-up and ignition stages, where heating by microwaves is believed to offer potential advantages over more traditional means.

Microwave heating and ignition of a combustible solid material is modeled and analyzed in the limit of small Biot number and large activation energy for a single-step, irreversible reaction. The model accounts for coupling between the microwave electric field and the temperature of the material in the context of a one-dimensional geometry and temperature-dependent material properties, including electrical conductivity, thermal diffusivity, and heat capacity. Since the period of a microwave is much less than the time required for heat to diffuse through a wavelength Maxwell's equations are averaged to yield the standard vector Helmholtz equation for the electric field. This is coupled to the diffusion equation for the temperature distribution of the sample, where the effects of electromagnetic heating are also averaged over the microwave period. Because the time required for heat to diffuse across the sample is much less than the time scale of heat loss through the sample boundaries the temperature distribution during the inert warm-up stage is almost spatially uniform. Perturbation methods in the high-activation energy limit yield local equations describing

the ignition process, which, as a consequence of the near-uniformity of the inert temperature distribution, can be of a more general form than those of standard ignition theory. Their numerical solution provides explicit ignition conditions, including the location of the point at which ignition occurs. It is found that, for a given material, the spatial distribution of hot-spots and onset of ignition can depend critically on the applied power of the microwave source.

Proposed developments include multi-dimensional effects and the microwave control of a developed reaction front. This work is being carried out in collaboration with John Bechtold, Gregory Kriegsmann, and Fu Li of New Jersey Institute of Technology.

JOHN A. BURNS

Applications of adaptive finite elements to sensitivity equations for optimal design

The computation of gradients for gradient based optimization may be accomplished through approximations of the continuous sensitivity equation. A typical procedure uses a simulation scheme to compute the solution of the forward problem and then linearizes around this numerical solution to obtain a sensitivity equation that can be solved by any number of numerical schemes. In many cases, this approach requires that functions representing higher order spatial derivatives be approximated by non-smooth piecewise defined elements. When used in a gradient based optimization scheme, such low order derivative approximations can produce large errors in the gradients and hence these approximate gradients are not useful in the optimization loop. One possible approach to this problem is to use high order finite elements to capture the spatial derivatives. This approach can eliminate the accuracy problem, but it also produces more complex systems and it requires more computational time and effort.

A second approach is to project the approximate derivatives onto smooth basis functions. This approach is similar to the ideas found in the adaptive finite elements literature where the goal is to construct local adaptive grids to minimize global error. Here we project the computed spatial derivatives for the forward problem onto a subspace of smooth basis functions to obtain an improved approximation to the sensitivity equation. This method has been applied to a simple 1D shape design problem and we are currently investigating variations of this method on other applications.

Future plans include the extension of these ideas to 2D flow optimization problems and to investigate the application of this idea to continuous adjoint methods. This work was done in collaboration with Dawn Stewart and Lisa Stanley.

INDRANEEL DAS

New algorithms for single and multiobjective nonlinear optimization

The main purpose of this study was the issue of engineering design under conflicting objectives, or what is commonly known as multicriteria or multiobjective optimization. An approach for finding many Pareto optimal points uniformly distributed in the Pareto set, which the author named Normal-Boundary Intersection (NBI), was already under investigation. This approach involved solving a nonlinear programming subproblem which, in general, would have both equality and inequality constraints. This motivated the development of an interior point trust region algorithm for the general nonlinear programming problem.

An interior point nonlinear programming algorithm was developed entirely while the author was at ICASE and a technical report containing the theoretical formulation and computational results was completed ["An interior point algorithm for the general nonlinear programming problem with trust region globalization"]. This was intended to be an extension of previous work by Coleman & Li and Vicente. Even though no theoretical proofs of convergence were provided, some computational results were presented which indicate that this algorithm holds promise. The computational experiments were geared towards improving the semi-local convergence of the algorithm; in particular high sensitivity of the speed of convergence with respect to the fraction of the trust region radius allowed for the normal step and with respect to the initial trust region radius were observed. The chief advantages of this algorithm over primal-dual interior point algorithms include better handling of the 'sticking problem' and a reduction in the number of variables by elimination of the multipliers of bound constraints. On the NBI front, the algorithm was refined and developed to a point where any number objective functions could be handled. Many new results with proofs were achieved, in particular the invariance of the obtained Pareto set with respect to function scaling and the relationships between NBI and other traditional methods. All of this can be found in the second technical report written by the author, entitled "Normal-Boundary Intersection: An alternate method for generating Pareto optimal points in multicriteria optimization problems".

Future plans include customizing the nonlinear programming algorithm developed for the NBI subproblem, perhaps using a multilevel step decomposition approach introduced by Natalia Alexandrov at LaRC. Some further refinements on NBI will be the topic of a forthcoming paper, including parallelizing the subproblem solutions.

The work on NBI was accomplished in collaboration with Dr. John Dennis of the Center for Research on Parallel Computation.

ADI DITKOWSKI

Error bounded algorithms for parabolic system

The SAT method for imposing boundary conditions has been applied in the past for the cases of diffusion and advection-diffusion equations, both in one and two dimensions. This algorithm allows one to use Cartesian mesh for "complex" domains.

A SAT algorithm was applied to the diffusion part of the compressible-viscous Navier-Stokes equations, as an example for parabolic system.

Future plans include extending this idea to hyperbolic and advection-diffusion systems in multi-dimension.

This work was conducted in collaboration with S. Abarbanel of Brown University.

PIERRE EMERIC

Application of smart material technology to damage detection in beam-like structures

Smart material technology is an active area of research, with promising applications that include, for instance, control systems and non-destructive evaluation. Among the different types of smart materials currently studied, structures with bonded piezoelectric ceramic patches are particularly interesting. When an electric field is applied, piezoelectric patches induce strains in the materials they are bonded to and, conversely, they produce a voltage when a deformation occurs in the material. As a consequence, these patches can act both as actuators and sensors, providing the host structure with smart material capabilities. The focus of this research is to apply smart material technology to the non-destructive detection and characterization of damage using vibration data from self-sensing/self-actuating beam-like structures.

We have developed mathematical models for structures such as thin plates and 2-D beams with crack-like defects and are currently analyzing computational methods for the detection and characterization of damage by detecting changes in the physical parameters such as stiffness, damping, mass density, etc. We have begun a series of experimental protocols designed to develop our abilities to verify our computational algorithms with experimental data. Further experiments will be conducted for the characterization of damaged structures.

This research was conducted in collaboration with H.T. Banks of North Carolina State University and W.P. Winfree of NASA Langley Research Center.

DANIELE FUNARO

Spectral elements for 3-D computations

The spectral element method, used for the approximation of boundary value problems, provides accurate computations, in comparison to solutions discretized by low-order finite-differences or finite-element methods, without encountering all the troubles that pure spectral methods exhibit (matrices with bad condition numbers, difficulties in dealing with complex geometries, etc.). This is why the interest in these techniques is growing.

I am currently working on the project of writing a 3-D code implementing spectral elements with algebraic polynomials of degree five (*pentacubes*), by collocation at Legendre type nodes. The main topic is the determination of approximated solutions to advection-diffusion equations in 3-D with small diffusive terms. The technique uses adaptive grids, depending locally on the differential operator and based on Legendre zeros, in order to stabilize the solutions. The research is in development and, at the moment, the study is concentrated on the search of very effective preconditioners for the iterative solution, by the multidomain approach, of the matrices related to the spectral element decomposition. The preconditioning matrices, acting on the interfaces unknowns, should take into account the direction and the magnitude of the transport terms. Several ideas are going to be tested in order to trim the code in an optimal way.

Part of the code has already been written and mainly deals with the organization of the vector of the unknowns, on the base of the multidomain subdivision, in order to address in the fast and proper way the exchange of informations at the interfaces (faces, sides and vertices). This part is not trivial if one wants to handle arbitrary geometries. The next step will be to construct the

solver in each single domain, and successively link the different local solutions through the iterative procedure.

JAMES GEER

Singular basis functions with built-in discontinuities

Series expansions of functions, such as Fourier series, perturbation series, etc., are often useful in developing numerical or semi-numerical, semi-analytical algorithms for the solution of differential equations. However, when only a partial sum of such a series is used, some "undesirable" effects (such as the Gibbs phenomena) may be present, or the partial sum may have "difficulty" approximating certain features of the solution, such as boundary layers, internal layers, or various discontinuities. The primary objective of this study is to completely eliminate Gibbs phenomena and its effects.

For a rather general class of functions f , the combination of a finite sum of certain singular basis functions with "built-in" singularities, along with a finite Fourier series, leads to a sequence of approximations which converges exponentially to f in the maximum norm, even though f may have a finite number of discontinuities. In particular, these approximations eliminate the Gibbs phenomena. In order to implement this approach, a knowledge of the locations and magnitudes of the jumps in f and its derivatives is necessary. A simple, accurate, and highly robust method of estimating these quantities from a finite number of the Fourier coefficients of f using a least squares technique has been developed and is being analyzed.

The class of basis functions with "built-in" singularities will be extended to include functions that have more general singularities, such as fractional power (e.g., square root) and logarithmic singularities. They will be applied to several model problems which either have discontinuities in the initial data and/or develop discontinuities (or "near-discontinuities") as time increases.

MIKHAIL M. GILINSKY

Advanced methods for acoustic, thrust and mixing benefits of nozzles, screws, propellers and fans

The theoretical and experimental research based on new methods for more efficiently mixing a supersonic jet with ambient air was continued. These methods are effective for both jet noise reduction and nozzle thrust augmentation. Several nozzle designs: Bluebell, Chisel and Screwdriver nozzles were tested theoretically. Numerical simulation was based on the Euler marching scheme of Krayko-Godunov. The top priority research is an application of a Bluebell nozzle concept for the two-contour nozzle with plug, which is the main nozzle design for the future hypersonic or supersonic aircraft engine. The experimental acoustic tests for this nozzle were started recently at NASA LaRC. Preliminary results are very promising. We made efforts to apply CFL3D and OVERFLOW NSE codes for numerical simulation of this flowfield. The comparison of different approaches will allow correct boundary condition formulation and choice of the best turbulence model.

A new concept for improving the working efficiency of propellers and screws was proposed in the invention disclosure NASA Technical Briefs published the note June 1996, which presents a short description of this invention. The concept is based on the Möbius strip-one-side surfaces. There

are several embodiments of such shapes. Each of them can be optimal for different application in different surrounding media.

Our research plan is to analyze the Möbius elliptical screw modification in the air flows. Definition of the optimal parameters provides maximum thrust, lift, mixing, or minimal noise as obtained by the analytical theory, numerical simulation and experimental tests. In the theory we will use a perfect gas model, Euler (potential and full) or Navier-Stokes equations with several turbulent models and NASA's package of CFD codes for main flow and acoustics flowfield numerical simulations. We will conduct similar research for some screw modifications in liquids: water, glycerine, viscous-plastic medium, which can be applied to enhance efficiency of the screws for boats and mixers. We will also test the proposed shapes in dispersive and dry media like sand, cement powders, coffee-beans and different mixtures.

STEPHEN GUATTERY

Lower bounds on the smallest nontrivial eigenvalues of Laplacians

The study of the connection between Laplacian spectra (particularly with respect to the second smallest eigenvalue λ_2 and its eigenvector) and properties of the associated graphs dates back to Fiedler's work in the 1970's. These properties have been used in graph algorithms e.g. for finding small separators; bounds on λ_2 are also useful in the analysis of spectral partitioning. λ_2 has been related to expansion properties of graphs, and can be used in determining if a graph is an expander. The Laplacian also often occurs in finite difference, finite element, and control volume representations of physical problems involving elliptic partial differential equations. These problems usually include a Dirichlet boundary condition that specifies that the values at a set of vertices are zero. To represent this in the Laplacian, the rows and columns corresponding to the boundary vertices are deleted. The resulting matrix is positive definite, and its smallest eigenvalue is of interest. Since the matrices are symmetric, their extreme eigenvalues can be used in computing their condition numbers, which are used in the study of iterative linear system solvers to estimate rates of convergence, and to analyze the quality of preconditioners. Thus the study of Laplacian spectra is of broad interest.

We introduced the path resistance method for lower bounds on the smallest nontrivial eigenvalue of the Laplacian matrix of a graph. The method is based on viewing the graph in terms of electrical circuits; it uses clique embeddings to produce lower bounds on λ_2 and star embeddings to produce lower bounds on the smallest Rayleigh quotient when there is a Dirichlet boundary condition. We showed that the best bounds that this method can produce for clique embeddings are the same as for a related method presented by Kahale that uses clique embeddings and edge lengths. Our method assigns priorities to paths in the embedding; we show that, for an unweighted tree T , the simple expedient of using uniform priorities for a clique embedding produces a lower bound on λ_2 that is off by at most an $O(\log \text{diameter}(T))$ factor. We have presented a useful tool that works well in various contexts, and that is particularly useful in analyzing classes of graphs in the study of asymptotic behavior.

Recent work has indicated that these techniques do not provide exact lower bounds because they do not represent the problem precisely. We hope to find a representation that will give an exact correspondence between embeddings and Laplacian eigenvalues.

This work was done in collaboration with Tom Leighton of MIT and Gary Miller of CMU, and will appear in the 8th ACM-SIAM Symposium on Discrete Algorithms.

GENE HOU

Improvement of simultaneous aerodynamic analysis and design optimization technique

The primary goal of this study is to develop the technique of Simultaneous Aerodynamic Analysis and Design Optimization (SAADO) that incorporates a validated computational fluid dynamics (CFD) code into an efficient aerodynamic design optimization procedure with minimal code modification. The current effort towards this goal focuses on the study of factors that affect the computational efficiency of SAADO.

SAADO incorporates design improvement within aerodynamic analysis iterations in order to gain computational efficiency for design optimization. The SAADO formulation treats both flow variables and shape variables as independent design variables. The SAADO solution process, however, introduces "inexact" sensitivity analysis to generate an optimization problem with shape variables being the only design variables. These "inexact" sensitivity equations and the flow equation are solved in a lock-step fashion to improve the design as well as the solution in each SAADO cycle. The feasibility study done previously demonstrated that for a turbulent transonic airfoil example, SAADO took the equivalent of 22 full aerodynamic analysis solution times to reach an optimal design. The current study has shown that the major factors which affect the computational efficiency of SAADO are: the efficiency and the method of solving "inexact" sensitivity equations; the number of function evaluations required by the line search algorithm; and, the accuracy of the flow and sensitivity variables in each SAADO cycle. The current study results in a revised SAADO formulation that uses "inexact" adjoint solution for sensitivity analysis and curtails the need of function evaluations for the line search. Numerical studies have not yet been completed to demonstrate the efficiency of the revised SAADO technique.

The immediate effort is to complete the numerical verification of the revised SAADO approach for a 2D example and then to extend it to 3D aerodynamic design optimization problems.

This research was conducted in collaboration with Clyde Gumbert and Perry Newman of the Multidisciplinary Optimization Branch, NASA Langley Research Center.

LELAND JAMESON

A wavelet-optimized, adaptive spectral method

Error control, grid adaptation, and order selection can all be accomplished quite naturally by using wavelet analysis. That is, wavelet analysis yields a set of coefficients indicating energy levels of flow variables as a function of scale and location throughout the domain. Such information is ideal for grid selection. Furthermore, wavelets analysis can distinguish between large smooth structures and small rough structures throughout domain, and such information is ideal for selecting the order of accuracy of differencing operators. Finally, by combining grid generation and order selection, one can keep the maximum error down to any prescribed user-selected tolerance in an efficient manner.

An adaptive numerical method has been created which adapts the numerical grid and the order of the differencing operator depending on the data. The numerical grid adaptation is at each level

of refinement a Chebyshev grid and this grid is refined locally based on wavelet analysis. Using the Chebyshev grid allows for very high orders of accuracy, 8,16, 20, etc., and in essence one has an adaptive order, adaptive grid, spectral method. The method has been named the Wavelet-Optimized Finite Difference Method2 or WOFD2.

The next step for the development of WOFD2 will be to extend to higher dimensions and to modify for complicated geometries.

R. MICHAEL LEWIS

Managing approximations and surrogates in optimization

The enormous computational cost of complex scientific and engineering simulations makes it impractical to rely exclusively on high-fidelity simulations for the purpose of design optimization. Instead, we will need to make as much use as possible of models of lower physical fidelity but lower computational cost, with only occasional recourse to expensive, high-fidelity simulations. Such use of approximation models is in keeping with engineering practice, where models of lower fidelity are widely used in preliminary design to explore the design space. The question that we are investigating is the systematic and algorithmic use of such approximations and surrogates in the context of optimization.

The approach we are pursuing is based on the trust region idea from nonlinear programming. The trust region mechanism provides a systematic response to both poor and incorrect prediction on the part of the approximation model, while not being so conservative as to retard progress when the approximation models do a good job of predicting the behavior of the "true" system. Furthermore, by monitoring how well the approximations are predicting the behavior of the system, the trust region mechanism suggest guidelines for changing or updating the approximation based on its predictive abilities. To date, we have determined the minimal requirements that approximations and surrogates must satisfy in order to be suitable for use in optimization algorithms, and have identified large classes of approximations that satisfy these conditions.

As part of this research, a week on approximation was organized this summer at ICASE, which brought together a group of roughly a dozen researchers from academia, industry, and NASA to discuss experience with using approximations and surrogates and to identify significant research areas in the field. A report of the week's discussion is in preparation.

Having settled the theoretical questions concerning the use of approximations, our efforts in this area will now focus on the much more interesting question of the development of practical algorithms. We intend to investigate the use of first-order kriging methods as general approximations, together with the use of local/global searches in optimization that attempt to take longer steps by using non-local models. We will also examine the use of problem-specific approximations, and the incorporation of information from models of many different levels of physical fidelity. We have identified two promising test cases, a structural optimization of a wing-box and the design model problem in the FIDO environment developed by the Langley MDO Branch. Two serious applications that we hope will become available to us in the next year are the optimization of helicopter rotorblades and the design of the Aerospike rocket engine nozzle for the Reusable Launch Vehicle.

This research was conducted in collaboration with N. Alexandrov of NASA Langley Research Center, J.E. Dennis, Jr. of Rice University, and V.J. Torczon of The College of William and Mary.

Pattern search methods for optimization

Pattern search methods remain popular with users of optimization because of their simplicity and surprising effectiveness in the unconstrained case. Our work in this area has focused on reducing the per-iteration cost of pattern search methods and extending this class of algorithms to general nonlinear programming problems.

Heretofore, to insure robustness of pattern search methods for unconstrained minimization, the minimal number of objective evaluations required in each iteration was $2n$, where n is the number of design variables. This made the cost comparable to a central finite-difference quasi-Newton method. By extending the class of pattern search methods and the attendant convergence analysis, either via the use of rank-order information or using the notion of positive linear dependence, we have reduced the per-iteration cost to as few as $n + 1$ objective evaluations, which is comparable to a forward finite-difference quasi-Newton method. Moreover, it is clear that no lesser amount of information will suffice to prove robustness, so our results are sharp.

We have also developed the first general class of robust pattern search methods for bound constrained minimization. It turns out that this is a proper subclass of the methods for unconstrained minimization. The restriction is necessary in order to avoid interactions between the directions of the pattern associated with the algorithm and the cone of descent directions for bound constrained problems. We have done some testing of these algorithms and found them to work quite well in practice; V.J. Torczon's collaborators at Rice and Boeing are currently developing an implementation to apply to a helicopter rotorblade application.

We intend to continue our work on generalizing pattern search methods to general nonlinearly constrained optimization. In the process of our work to date, we have developed a fairly precise understanding of the types of geometrical and analytical problems we must guard against if we are to develop provably robust algorithms. The key is to insure that the patterns have a sufficient density of points relative to the tangent cone of the constrained region. The main obstacle is the absence of Lagrange multiplier estimates since pattern search methods are derivative-free methods, and to avoid the problems that bedevil steepest descent in the non-smooth case, if we use exact penalty functions.

This research was conducted in collaboration with V.J. Torczon of The College of William and Mary.

Qualitative nature of Hessians for problems associated with PDE

In optimization of systems governed by partial differential equations, the nature of the governing PDE affects the nature of the Hessian of the objective and constraints. At the same time, accurate approximation of second derivatives greatly improves the efficacy of nonlinear programming algorithms. This research focuses on examining the relationship between the nature of the governing PDE and the nature of the Hessian.

We have developed an approach to sensitivity calculations for systems governed by PDE that reveals some of the relationship between second derivatives and the nature of the governing PDE. This perspective is based on the connection between first and second derivatives in such problems to the reduced gradient and reduced Hessian. In particular, we can see how the solution operator for the PDE enters into the structure of the Hessian operator.

The structure of the Hessian is such that if the governing PDE is elliptic or parabolic, the associated Hessian will be something like a pseudodifferential operator, while if the governing PDE is hyperbolic, the Hessian will be something like a Fourier integral operator. In particular, the Hessian associated with hyperbolic problems may propagate singularities. At the same time, we can show that many Hessian approximations, such as the Gauss-Newton approximation for least-square problems, will not propagate singularities and thus as qualitatively quite wrong.

We intend to pursue the analytical aspects of the relationship between Hessians and PDE, and to explore the practical, numerical consequences of this relationship.

JOSIP LONCARIC

Spatial structure of separation feedback control

Flow separation over wings at high angles of attack limits lift, which is particularly important at takeoffs and landings. Low limit on lift necessitates larger and heavier wings, which has a disproportionately large effect on the useful aircraft payload. Several techniques (heating, blowing, sound) are known to delay flow separation, with open-loop unsteady inputs being more effective than steady inputs. New technologies (e.g., MEMS) hold the promise of constructing smart aircraft systems capable of using feedback control to maintain the desired flow state. While much *can* be done to the flow, the question of what *should* be done remains open. Traditionally, control design is carried out after the system is already specified. We intend to compute optimal feedback gains not with the goal of implementing them, but with the goal of providing new insight helpful in designing smart high lift systems.

This project will compute the minimum effort distributed flow control feedback operator and investigate its spatial structure. Those spatial regions where feedback sensitivity and action are large will guide placement of a finite number of sensors and actuators, suitable for control implementation later. Our approach begins with the recent work on optimal feedback control of incompressible Navier-Stokes equations. The main remaining obstacle is computational cost, which we aim to reduce by using spectral flow representation on an efficient grid. The laminar 2D problem is expected to be both physically reasonable and computationally feasible. We are constructing the nominal flow state by combining the exterior potential flow with an approximately steady boundary layer, so that a standard LQR problem may be obtained by discretization and linearization of the Stokes operator and the inertia term.

The problem is conceptually defined and the numerical algorithm is largely worked out. We intend to refine the selection of the nominal flow and the grid, and then proceed with the coding. Preliminary runs will be done at coarse resolution (16^2 fully de-aliased), and the spatial structure of these feedback operators will be used to tune the grid for the final runs (64^2 fully de-aliased). These distributed control results will enable us to guide further work in boundary control. This numerical study is also expected to lead to useful analytical approximations in the feedback control of boundary layers.

This research was conducted in collaboration with Brian Allan of ICASE.

Placement of sensors and actuators in noise control

Interior noise on aircraft is largely produced by vibration of fuselage panels in response to forcing by exterior pressure fluctuations. Passive damping of such vibration exacts a significant

weight penalty, which can be alleviated by active control. Successful active noise control strategies have been demonstrated by attaching piezoelectric patches to the panel surface, but the question of optimal placement of sensors and actuators remains unanswered. Motivated by this problem, we aim to develop a model-based optimal control theory approach to the design of such distributed parameter systems.

Structural vibration of panels and interior acoustic radiation can be modeled by differential operators which belong to the class of spectral systems. If one could measure everything and act everywhere, what should one do? In this thought experiment the control input is allowed to act everywhere, which allows simultaneous diagonalization of both system dynamics and control. The infinite dimensional LQR problem therefore decouples into an infinite series of low dimensional LQR problems which are easy to solve, each of which contributes a term to the optimal feedback operator K . Spatial regions where K exhibits large sensitivity or action suggest good locations for sensors and actuators. We have obtained analytic expressions for K (in the distribution sense) for a simply supported plate and several control objectives. The results indicate that positional stabilization requires control of only the first few modes, but that minimizing the plate velocity, acceleration or energy requires high bandwidth. Furthermore, high frequency behavior of the optimal feedback gains at low internal damping is a singular perturbation of the undamped case.

Human ears are sensitive only to certain range of frequencies which means that the required control bandwidth is also limited. We intend to explore the spatial structure of optimal K corresponding to bandwidth limited objective functions, and then extend the results to the clamped plate case through numerical experiments. We also intend to explore coupling of structural and acoustic modes, as well as different input disturbances.

DIMITRI J. MAVRIPLIS

Adaptive meshing for mixed element unstructured meshes

The use of hexahedral and prismatic elements in unstructured meshes in addition to traditional tetrahedral elements offers the possibility of increased accuracy and reduced overheads. In previous work, a discretization and solution strategy for mixed-element unstructured meshes was developed, and gains in efficiency due to the use of different element types were demonstrated. The present work seeks to develop adaptive meshing techniques for mixed element unstructured meshes. The ability to adaptively refine mixed element meshes is important since adaptive meshing represents one of the principal advantages of unstructured meshes.

Mesh refinement is achieved by element subdivision. For each element type, the permissible subdivision patterns (isotropic and anisotropic subdivision) are defined, classified and incorporated into a generic element subdivision library. The hierarchical subdivision history is also stored. This enables de-refinement by retracing the subdivision history. More importantly, the hierarchical information is utilized to avoid multiple levels of anisotropic refinements. For anisotropically refined cells (i.e. which occur at interfaces between refined and unrefined regions) any further refinement requires the deletion of the current anisotropic children cells, and the isotropic refinement of the corresponding parent cells. This is necessary to avoid degradation of mesh element quality with additional refinement levels. Adaptive refinement of fully tetrahedral meshes results in fully tetrahedral refined meshes, while refinement of hexahedral meshes or combined hexahedral, pyramidal,

prismatic and tetrahedral meshes results in refined meshes contain combinations of these types of elements.

The refinement routines are to be coupled with the flow solution and multigrid schemes in order to compute flows in adapted grids of mixed elements. Increases in accuracy and efficiency over non-adapted meshes and fully tetrahedral meshes will be demonstrated on large cases of the order of 5 million grid points.

Unstructured multigrid convergence acceleration for highly stretched meshes

For high-Reynolds number viscous flow simulations, efficient resolution of the thin boundary layer and wake regions requires mesh spacings several orders of magnitude larger in the normal direction than in the streamwise direction. This extreme grid stretching results in poor multigrid convergence rates, usually one to two orders of magnitude slower than those observed for equivalent inviscid flow problems without grid stretching. The purpose of this work is to devise improved multigrid techniques for acceleration convergence on anisotropic grids of this type.

The approach taken consists of using implicit line solvers in the direction normal to the grid stretching combined with semi-coarsening or directional coarsening multigrid to alleviate the stiffness due to grid anisotropy. A graph algorithm has been implemented which combines edges of the original mesh into lines which follow the direction of maximum coupling in the unstructured mesh. A similar algorithm is used to selectively coarsen the original fine mesh by removing points along the directions of strong coupling, thus recursively generating a hierarchical set of coarse meshes for the multigrid algorithm. Using this approach, convergence rates for viscous flows which are independent of the degree of mesh stretching can be obtained. These multigrid convergence rates are close to those achieved on equivalent inviscid flow problems with no mesh stretching.

Although present results are encouraging, they have only been demonstrated for simple two-dimensional flows. The methodology is to be implemented in three dimensions and combined with better preconditioning techniques to further improve efficiency.

LUCA F. PAVARINO

Parallel preconditioners for mixed spectral element methods for elasticity and Stokes problems in three dimensions

The pure displacement formulation of the linear elasticity equations can suffer from the phenomenon of locking when the Poisson ratio ν tends to $1/2$ (almost incompressible case) and standard finite element discretizations are employed. This problem can be overcome by using a mixed finite element formulation. The resulting problem is equivalent to the steady Stokes system in case of homogeneous Dirichlet boundary conditions on the whole boundary of the domain. Therefore, discretization techniques used for incompressible flows can be adapted to almost incompressible elasticity problems. Spectral element discretizations, which have been successfully employed for large Stokes and Navier-Stokes 3-D calculations, can be easily adapted to elasticity problems. The discrete system obtained is larger and indefinite and the inf-sup condition has to be satisfied by carefully choosing the finite element spaces. Good parallel preconditioners for such systems are the key to the practical solution of large 3-D elasticity and Stokes problems. The objective of this research is to introduce and analyze parallel preconditioned iterative methods with convergence rate independent of the Poisson ratio ν and therefore well suited for incompressible flows or materials.

We have studied and implemented two spectral element methods for the mixed formulation of the Stokes system and linear elasticity in 3-D. The first of these methods is analogous to the $Q_n - Q_{n-2}$ spectral element method of Maday, Patera and Ronquist, while the second is known as $Q_n - P_{n-1}$. A basis for the spectral spaces is built using Gauss-Lobatto-Legendre (GLL) quadrature. We have considered two strategies for the iterative solution of the resulting saddle point problem: a) solve the whole indefinite system with block-diagonal or block-triangular preconditioners and b) eliminate the interior variables and solve the resulting Schur complement by a substructuring domain decomposition method of wire basket type. In case a), domain decomposition techniques could be employed as preconditioners for the individual blocks. Using recent results by Klawonn, we have proved and shown numerically that the convergence rate of the proposed methods is independent of ν and the number of spectral elements N , and only mildly dependent on the spectral degree n via the inf-sup constant. In case b), the substructuring method seems to have a similar convergence rate, but more research is needed to fully understand it.

Future work will include overlapping domain decomposition methods and preconditioners based on h -version finite elements or finite differences on the GLL mesh. All these techniques should then be extended to the Navier-Stokes equations.

Part of this research was conducted in collaboration with Olof B. Widlund of the Courant Institute of Mathematical Sciences (NYU).

PHILIP L. ROE

Adaptive meshes derived from minimization principles

The computational solution of any problem requires a discrete grid on which the solution will be defined. The choice of grid may strongly influence the quality of the solution. An effort has been initiated to generate such grids in an optimal way.

Early attempts to define grid quality did so in terms of orthogonality, skewness and smoothness. It has become clear that none of these properties is desirable in itself, but that the grid needs rather to reflect properties of the governing equations, the specific solution of those equations, and possibly also of the discrete solver. A least square minimization principle that optimizes grid placement and connectivity has been formulated. Crucial to success is the norm in which the minimization is carried out. When this is specified correctly, the resulting grids have striking properties for simple model problems of both elliptic and hyperbolic type. A paper describing the early stages of this research has been written for the proceedings of the ICASE/LaRC Workshop on Barriers and Challenges in Computational Fluid Dynamics.

The immediate future objectives are to extend the strategy to nonlinear problems with discontinuities and to three dimensions.

CHI-WANG SHU

Discontinuous Galerkin method and TVD Runge-Kutta methods

Our objective is to study and apply high order finite difference (ENO type schemes), finite elements (discontinuous Galerkin) and spectral methods for problems containing shocks. This will enable us to capture complicated solution structure over long period of time (such as the problems in aeroacoustics) with a relatively coarse grid.

The investigation of the discontinuous Galerkin finite element method, which is carried out jointly with Harold Atkins at NASA Langley, is in the phase of studying the discretization of viscous terms. Direct application of discontinuous Galerkin method leads to an inconsistent result. We are investigating the application of discontinuous Galerkin method to the first order system by introducing additional variables. The implemented code is being tested for acoustic related problems. Jointly with Sigal Gottlieb of Brown University, we are investigating the optimal TVD Runge-Kutta time discretizations of second, third, and fourth orders, and those with low storage requirements.

Research will be continued for high order methods in finite difference, finite elements and spectral schemes.

DAVID SIDILKOVER

A simple essentially optimal multigrid solver for the inviscid flow equations

The main objective of this work is to develop a simple essentially optimal multigrid solver for the steady Euler equations (both incompressible and compressible), i.e. a solver whose efficiency is similar to that of a solver for Poisson or Full-Potential equations.

A simple discretization of the Euler equations (both compressible and incompressible) that facilitates such a solver was constructed recently. The discrete equations to be solved at each node are "assembled" from the residuals of the Euler system on the grid elements having this node as a common vertex. The elliptic factor separates in the form of Poisson (incompressible case) or Full-Potential (compressible case) operators acting on the pressure. It was demonstrated that the optimal multigrid efficiency can be obtained for some model problems using this approach.

The current work concentrates on implementation of the algorithm in the framework of some engineering codes, capable of treating complex geometries. This is being done both in the context of unstructured triangular and structured Cartesian grids. The essentially optimal multigrid efficiency was already demonstrated for incompressible Euler equations in some complex geometry cases. The near future plans include addressing of some even more realistic test cases and generalizing the implementation for incompressible laminar Navier-Stokes and compressible subsonic Euler equations.

This is a joint work with Drs. T.W. Roberts and R.C. Swanson of NASA Langley Research Center.

Multidimensional upwinding and textbook multigrid efficiency

A genuinely multidimensional high-resolution discretization for the steady compressible Euler equations was constructed recently. One of the fundamental advantages of this approach is that the genuinely multidimensional high-resolution mechanism (unlike the standard one) does not damage the stability properties of the scheme. The next main objective is to construct an optimally efficient multigrid solver that relies on the genuinely multidimensional upwind discretizations. The main advantage of this approach is its generality – it is expected to be extendable to the cases/regimes other than steady inviscid subsonic flow.

It is well known that for the advection dominated problems the coarse grid provides only a fraction of needed correction for certain error components. A known cure for this problem is to

separate the treatment of the advection part from the rest of the system. A simple way to achieve this is through the separate discretization of different factors. A more general approach is to discretize the equations in some primitive form but to design a special relaxation (of the distributive Gauss-Seidel type) that distinguishes between the different co-factors of the system. A necessary condition for this, however, is that the discretization must be *factorizable*, i.e., the determinant of the corresponding matrix of discrete operators should be a product of approximations to different co-factors of the system.

A recent study has revealed that some genuinely multidimensional schemes are factorizable. The current work concentrates on exploring the possibility to obtain the textbook multigrid efficiency in subsonic regime using the genuinely multidimensional upwind schemes, in particular, on the design of the so-called distributive Gauss-Seidel relaxation.

RALPH SMITH

Galerkin approximation of thin shell dynamics

A large number of models for physical phenomena incorporate shell-like dynamics. Broadly stated, such models are employed for structures in which the thickness is significantly smaller than orthogonal dimensions, and coordinate motions are potentially coupled due to geometrical effects such as curvature. Since analytic solutions are typically unavailable, efficient approximation techniques must be developed to permit implementation of the models in simulation studies, parameter estimation routines and control applications. The use of finite difference techniques is severely hampered by the high order of the differential operators and boundary conditions. Compatibility and "locking" phenomena must be accommodated when considering finite element methods. The objective in this research is to develop an efficient Galerkin method for approximating thin shell dynamics which is suitable for control applications.

A Galerkin method for approximating cylindrical shell dynamics has been developed and numerically tested. A spline basis is employed in the axial direction and Fourier expansions are used in the periodic circumferential direction. This approximation framework is suitable for shells containing nonhomogeneities due to smart material actuators/sensors and provides sufficient flexibility to encompass a variety of boundary conditions found in experiments. On its own, this Galerkin method provides an accurate means of approximating both shell dynamics and static components such as natural frequency and mode shapes. The method is also sufficiently flexible to permit direct linking with acoustic codes so as to provide an efficient framework for approximating coupled structural acoustic dynamics.

We are currently melding this approximation method with a finite dimensional LQR theory to provide a means of simulating full state feedback control of shell dynamics using piezoceramic actuators. Included in this latter component are investigations regarding optimal patch number and placement to yield maximum attenuation for various external excitations to the shell.

This work was done in collaboration with Ricardo del Rosario of Iowa State University.

Compensator design employing direct integration methods to obtain adjoint solutions

The application of optimal control theory to systems with exogenous disturbances yields a feedback law comprised of state estimates and solutions to a corresponding adjoint or tracking

equation. For many applications, it is reasonable to assume that the exogenous disturbance is composed of an underlying combination of periodic signals with white noise superimposed. The objective in this investigation is to utilize the underlying periodicity in the disturbance to develop direct integration methods for the adjoint equations which can be experimentally implemented.

In collaboration with David Cox, Flight Dynamics and Controls Division, an efficient and stable technique for integrating the adjoint equations has been developed. The method is based on the assumption that the underlying disturbance consists of a superposition of temporally periodic signals. An explicit adjoint solution is then obtained through a variation of parameters formulation based on the eigendecomposition of the system matrix. Phase, frequency and magnitude values are obtained through nonlinear filtering techniques. In this manner, the calculation of the adjoint solution at any time is reduced to matrix multiplication. This eliminates the necessity of numerically integrating the adjoint system from a future time to the present to obtain a current solution.

Numerical simulations have demonstrated that the direct integration method is extremely efficient when applied to structural systems modeling experimental apparatuses. Current efforts are directed toward the experimental implementation of this technique. This involves the design of circuits and filters to obtain the phase, frequency and magnitude of the disturbance, and the programming of data acquisition systems required for implementation.

This work was done in collaboration with H.T. Banks of North Carolina State University.

SHLOMO TA'ASAN

Optimal multigrid solvers for inviscid flows

Recent observations regarding the decomposition of the Euler equations into elliptic and hyperbolic parts have led us to believe that much more efficient solvers can be developed. The present performance of Runge-Kutta based methods together with coarse grid acceleration is limited to convergence rate of .75 per multigrid cycle, and this is observed in the numerical simulations. The objective of this research is to further develop ideas that have shown Poisson like multigrid efficiency for the subsonic Euler equations.

The main difficulty in reaching Poisson-like efficiency for multigrid algorithm for the Euler equations is related to the treatment of hyperbolic parts of the system. These are related in the subsonic case to the propagation of entropy and total enthalpy along streamlines. In the supersonic case all variables propagate in a certain way. Multigrid techniques have found to be not efficient enough for accelerating propagation problems. This implies the use of marching techniques in the direction of the flow for certain components of the system (entropy, total enthalpy, and in 3D also helicity). The conjecture is that the proper treatment of the different subsystems (elliptic and hyperbolic) will lead to optimal performance. The purpose of this research is to find the right implementation of these ideas for all flow regimes.

A recent attempt to use nonstaggered grids led to a formulation in terms of vector potential in 3D. This formulation allows the use of vertex based schemes which are symmetric discretizations of the elliptic parts of the system, thus reducing the artificial viscosity. The hyperbolic parts use upwind discretization. While marching techniques are used for the entropy, total enthalpy and helicity, an efficient smoother is constructed for the elliptic parts. This relaxation is then accelerated using coarse grids in the usual way.

The method has been demonstrated for the incompressible case, both in 2D interior and exterior flows, and in 3D for interior flows. Poisson-like efficiency have been obtained.

The near future plan is to focus on 3D exterior problems, where a flow around a wing will be demonstrated. The step following this will involve the compressible 3D equations for exterior flows.

ELI TURKEL

Preconditioning

We are testing preconditioning for three dimensional flows with complicated geometries. It is well known that most codes have great difficulties solving problems when the Mach number is too low. These difficulties involve both very slow convergence rates to a steady state and also accuracy difficulties even once a steady state is achieved. Though the flow is almost incompressible there are numerous instances that one cannot use an incompressible code e.g. transonic flow in a small portion of the domain. Preconditioning is a technique for overcoming these deficiencies. This is based on changing the eigenvalues of the inviscid system so as to more equalize them for low speed flow. At the same time the artificial viscosity, or Roe matrix, is modified. These two changes improve both the convergence rates and the accuracy.

We are calculating the inviscid flux terms as in the standard TLNS3d approach. This is then multiplied by a matrix which simultaneously preconditions the residual and converts it to (p, u, v, w, T) variables. We then add the artificial dissipation and calculate the primitive variables at the next stage of the Runge-Kutta method. Finally the conserved quantities are calculated. Thus the preconditioned quantities are used for both residual smoothing and multigrid transfers. We have also developed a matrix artificial viscosity that is appropriate for the preconditioning. This has been tested on many different complex configurations for both external and internal flows.

Other techniques for accelerating the convergence to a steady state are being developed. These include Jacobi preconditioning and including aspect ratio information in many of the parameters of the algorithm.

This is joint work with Veer Vatsa of NASA Langley Research Center.

Cusp & slip schemes

The second project is a study of the CUSP scheme devised by Antony Jameson. This scheme uses a splitting of the convective and pressure terms and gives an alternative artificial viscosity to the matrix viscosity previously developed by Swanson and Turkel. We are also studying various limiters including SLIP.

We have programmed up the HCUSP and SLIP and done numerous comparisons for subsonic, transonic and supersonic flow for both the inviscid and viscous equations. We have also combined the CUSP scheme with preconditioning for low speed flow. In general CUSP gives very sharp shocks but the global quantities such as drag and lift are less accuracy with CUSP than with the matrix valued artificial viscosity. The computer time for both are approximately the same. Both schemes are much more accurate than the original scalar artificial viscosity.

Our plan is to work on schemes that combine the best advantages of all these algorithms.

This is joint work with C. Swanson of NASA Langley Research Center and R. Radespiel of DLR.

BRAM VAN LEER

Local preconditioning for the Euler and Navier-Stokes equations

Local preconditioning matrices for the Euler and Navier-Stokes equations have been shown to yield significant single- and multi-grid convergence acceleration. The reason these are not yet in widespread use is their lack of robustness, especially when stagnation points are present in the flow solution. It has been pointed out by Darmofal and Schmid (1995) that the lack of convergence or even instability near stagnation points is caused by certain eigenvector pairs, associated with distinct waves (governed by the preconditioned equations) that become parallel at vanishing Mach number. A “fix” is to bound the Mach number away from zero, but the cut-off is rather arbitrary, and the construction of the proper artificial-viscosity terms becomes more cumbersome.

A matrix proposed by Dohyung Lee, doctoral student at the University of Michigan, has been studied which seems to have a healthier eigenvector structure for $M \rightarrow 0$, although not necessarily for $M \rightarrow 1$. This matrix can be related to the earlier known matrices of van Leer (symmetric) and Turkel (triangular); in fact, it has been modified to fit into a one-parameter family with the other two. The modification is negligible at $M = 0$, which is where the matrix would be used anyway. When M grows from 0 to 1, the general matrix changes from Lee’s via Turkel’s into van Leer’s, which has a good eigenvector structure for $M = 1$. In one dimension it is possible to create two versions of Lee’s matrix that have exactly orthogonal eigenvectors for *all* Mach numbers. This property is lost in two dimensions.

Further search among optimal and maybe slightly suboptimal preconditioning matrices (resulting condition number < 3), in order to find still more closely orthogonal eigenvector sets, and that uniformly over the entire Mach-number range will be conducted. This should immediately translate into greater numerical robustness in practice.

HONG ZHANG

Developing parallel algorithms for multiple objective linear programs

A fundamental problem in multiple objective linear programming is to determine the efficient set that often has a large or infinite number of solutions. The problem becomes even more complex if a multiple objective linear program (MOLP) is of a large size. Some well established sequential algorithms for MOLPs, such as the software package ADBASE, are extensions of the simplex method for linear programming. The growing importance and enormous computational expense of MOLPs motivated the development of parallel algorithms capable of solving very large size problems.

Based on ADBASE, a parallel ADBASE algorithm, that significantly accelerated the solution process of MOLPs on multiprocessors, was developed jointly with M.M. Wiecek at Clemson University. However, the original ADBASE bookkeeping strategy of maintaining a large unsorted list of efficient solutions was found to be a major bottleneck for sequential processing, as well as for parallel processing. Our current work addresses this issue by comparing three representative bookkeeping strategies, selecting the optimal strategy, incorporating it into ADBASE and the parallel ADBASE algorithm, and implementing the algorithms.

Future plans include finding a reliable and efficient way of replacing the global bookkeeping with local lists, and applying the algorithms to real decision-making problems.

FLUID MECHANICS

FRANCOISE BATAILLE

Numerical simulation of compressible turbulence

We investigate compressibility effects on turbulence in homogeneous and isotropic flow. Using EDQNM (Eddy Damped Quasi Normal Markovian) as a two point closure in Fourier space, simulations of forced isotropic turbulence lead to the conclusion that the compressible (i.e., irrotational) velocity spectrum scales as $k^{-5/3}$ in the inertial zone for short times and has a long time asymptotic state which scales as $k^{-11/3}$. The aim of the study was to find out whether this behavior could be reproduced through a large-eddy simulation of isotropic turbulence. We tested several initial conditions. Among these, we considered both analytical (with a $-5/3$ slope) and experimental (the spectrum of Comte-Bellot/Corrsin) initial spectra. Furthermore, the effect of Reynolds number was investigated. The $k^{-11/3}$ behavior could be recovered, even on grids as coarse as 64^3 . Both the Smagorinsky and the Dynamic subgrid scale models have been implemented. The same tests performed with these two models led to similar results. Therefore, it does not appear that a more complex subgrid-scale model is required for the calculation of spectral slopes in isotropic turbulence.

Turbulent boundary layer with blowing

This study involves a turbulent boundary layer problem linked to a fluid transfer (blowing) coming from a porous plate. Numerical investigations, using Fluent, showed that blowing increases the boundary layer thickness and decreases friction factor. These results have been obtained (at the INSA-Lyon in France) using a new model of the injection which leads to a very good agreement with experimental data.

In collaboration with B. Younis, an attempt was made to duplicate the results using another code.

Discussion with C. Streett of NASA Langley Research Center led to the conclusion that, in the future, a Large Eddy Simulation of a turbulent boundary layer with blowing (using the new injection model) would be very interesting. It could be done by modifications of his large eddy simulation code.

ALVIN BAYLISS

Numerical simulation of jet acoustics and structural/flow/acoustics interaction

The objective of this research is to simulate the generation of sound in a jet together with the response and radiation of nearby flexible aircraft type panels excited by jet sound. In practice aircraft are composed of flexible panels which are excited by sound from the jet. This loading causes panel vibration which can lead to increased interior noise levels and structural fatigue. In order to be able to control panel vibrations it is necessary to simulate panel response to consistent loading (i.e., loading representative of sound emanating from a jet).

We simulate this problem using a fully coupled approach in which the equations governing the jet and associated acoustics are fully coupled to equations describing the panel motion. The

equations are solved together with the pressure computed from the fluid used as a forcing term for the panels, and the velocity obtained from the panel dynamics employed as a boundary condition for the fluid dynamical computation. Our computation includes the effect of inviscid instability waves in the jet as well as disturbances generated from the nozzle lip. To date we have identified several different regimes of jet response. In particular, we have completed research on high subsonic jets. We have found that the response of the jet is dominated by a fundamental frequency and harmonics and this significantly effects the panel response and radiation. A paper describing this work will appear in the AIAA Journal. We have also studied this problem for supersonic jets. We have found that radiation along the Mach angle leads to a significant sensitivity in panel response depending on panel location. A preliminary report on this work will be presented at the AIAA Aerospace Sciences meeting.

Future plans involve extending our simulation to account for axisymmetric and nonaxisymmetric effects for circular jets. An extension of the computer program to accomplish this is under way.

Collaborators in this research are Lucio Maestrello of NASA Langley Research Center and Charles C. Fenno, Jr. of the National Research Council.

AYODEJI DEMUREN

Near-wall turbulence model based on RNG theory

Turbulence models are usually derived on the assumption of high Reynolds numbers, which enables significant simplification and the neglect of direct viscous effects. However, in the immediate vicinity of walls, viscous effects are important and Reynolds numbers are low in the viscous sublayer and the buffer layer. The turbulence models must be modified in this region, but most existing modifications are ad-hoc and contain little physical reasoning. The objective of the present study is to use Renormalization Group Theory (RNG) to develop low-Reynolds number turbulence models which have a sounder basis.

High-Reynolds number $k-\epsilon$ models can be derived by assuming the Kolmogorov spectrum based on the $-5/3$ law. In low-Reynolds number turbulence, the spectrum deviates from this. The approach then is to use a modified spectrum to derive low-Reynolds number effects on the turbulence. The derivation is completed, and model computations are being tested against DNS and LES results of channel flow. Initial results are promising, but we see a need to simplify the equations.

The model will be simplified and then tested in separated flows as a precursor to application in high-lift configuration.

SHARATH S. GIRIMAJI

Algebraic Reynolds stress model for curved flows

Curved turbulent flows continues to be a major challenge at the two-equation level of turbulence modeling. A simple, robust turbulence model that adequately accounts for the curvature effects without violating basic invariance principles is the objective of this research.

Further improvements to the previously developed algebraic Reynolds stress model for curved flows were made. This includes the recognition that the weak-equilibrium assumption indeed should produce the long-term steady-state solution of the Reynolds stress closure equations.

Validation of the model and modeling assumptions is underway using the DNA and LES data of a curved circular pipe flow.

This research was conducted in collaboration with B.J. Boersma of Delft University of Technology.

Algebraic modeling of scalar flux

Turbulence has a profound effect on heat and mass transfer. An algebraic model for passive scalar and thermal flux would be very useful in computing turbulent combustion and convection.

Models for turbulent passive scalar flux and thermal flux in mixed and natural convection are being developed using the algebraic stress modeling methodology. The evolution equation of the scalar flux is subject to the weak equilibrium assumption and representation theory is invoked to accomplish the model development.

Validation of the ensuing models using DNS data is the next step of this project.

Non-equilibrium algebraic Reynolds stress modeling

Algebraic Reynolds stress models work best at or very close to equilibrium turbulence. In many practical flows, significant departures from equilibrium turbulence are encountered. Currently, full Reynolds stress closure models are required to capture these transient effects. An algebraic model that can account for moderate departures from the equilibrium state will be a valuable engineering tool, and such is the objective of this project.

The possibility of obtaining an algebraic model when some of the equilibrium assumptions are relaxed is being explored. It is not completely clear if such a model can be obtained starting from the Reynolds stress closure equation. However, there are some promising indications that justify further investigations.

M. EHTESHAM HAYDER

Prediction of jet noise

This study focuses on utilizing the results from advanced turbulence closure model predictions of supersonic jet flow-fields in the computation of the associated aerodynamic sound fields.

We compute the mean turbulent flow-field by solving the Reynolds averaged Navier Stokes equations. The Lighthill acoustic analogy is used to compute the far-field sound from results of a computed jet. In the acoustic formulation, the source strength is expressed as a fourth-order velocity correlation. This correlation can be approximated as a linear combination of second-order correlations. We use individual stress components from computed field and thus do not use the isotropy assumption. In an alternate approach, we solve the linearized Euler equations with a modeled noise source term. This source term is determined from the computed turbulent mean flow field. We tried a stochastic technique to determine the noise source.

We plan to improve our noise source model and compute radiated noise field.

This research was conducted in collaboration with Jack Seiner of NASA Langley Research Center, Bob Rubinstein, and Y. Zhou of ICASE.

JEFFREY HITTINGER

Computational methods for sound-generating flows

As propulsion systems on aircraft have become quieter, airframe noise has become a more dominant contributor to the overall noise spectrum of aircraft. Since airframe noise is often characterized by three-dimensional, separated, vortical flow, numerical simulation of the acoustic near field is prohibitively expensive using current algorithms for the unsteady, compressible Navier-Stokes equations. The objective of this research is to develop more efficient numerical procedures for computation of the sound-generating flow region of airframe noise sources.

During my first two months at ICASE, a great deal of time was spent familiarizing myself with the subject of airframe noise by reviewing the literature and conferring with NASA and ICASE scientists working in related areas. From this, the significance of airframe noise and the difficulties with its calculation became apparent. Current Navier-Stokes algorithms for high-Reynolds number flow are extremely expensive; implicit formulations are used because of the unreasonable numerical stability constraints imposed on explicit schemes. Ideally, we want to develop a local, explicit scheme to solve viscous flow problems for which the time step is restricted by accuracy rather than stability. In addition, this new method would lend itself to efficient grid adaptation to reduce the cost of the calculation. Thought along these lines lead to the idea of solving the 10-moment equations instead of the Navier-Stokes equations. While this might appear to be a complication of the problem, recent work suggests that the mathematical nature of the 10-moment equations would lend itself to more efficient computation. It is also suspected that this approach could be easily coupled to an adaptive mesh refinement strategy. Currently, analysis is being conducted to justify the feasibility of the 10-moment-equations approach.

If the results of this feasibility analysis are favorable, research will proceed with the implementation of the method proposed above. Several issues such as boundary conditions, turbulence modeling, and parallelization will have to be addressed. Test cases will be run to demonstrate the efficiency of the method in comparison with compressible Navier-Stokes algorithms.

This work is being conducted under the supervision of Philip Roe of The University of Michigan.

FANG Q. HU

Computational aeroacoustics and absorbing boundary conditions

The objective of this work is to develop highly accurate numerical schemes and boundary conditions suitable for Computational Aero Acoustics (CAA).

For high quality numerical solutions of aeroacoustics problems, accurate non-reflecting boundary conditions are of critical importance. A new technique has been under development using the Perfectly Matched Layer equations recently proposed for the Euler equations. In the past six months, research efforts have been directed in these areas : (1) non-linear simulations, (2) inflow boundary conditions and (3) CAA Benchmark problems. In collaboration with M.E. Hayder (ICASE) and M.Y. Hussaini (Florida State), application of the PML technique to non-linear simulations of plane shear layers and axisymmetric jets have been carried out. Assuming that the actual mean flow was unknown, the PML equations were constructed using the concept of a pseudo-mean flow for non-linear simulations. The preliminary results are very encouraging. Moreover, a methodology of

using PML absorbing boundary condition as an inflow boundary condition, in which the incoming waves, such as a plane acoustic or vorticity wave, can be specified in the inflow region while the out-going waves will be absorbed simultaneously, has been proposed. It consists in separating the incoming and out-going waves and applying the absorbing boundary condition on the out-going waves. Numerical results have been satisfactory. Finally, the PML technique has been applied to CAA benchmark problems including the acoustic scattering by a cylinder, duct acoustics and turbomachinery noise. The numerical results will be presented in the Second Workshop of CAA, Nov. 1996.

The future work will further extend the PML technique and its analysis to axisymmetric and 3D simulations and the Navier-Stokes equations, with the goal of computing noise directly in realistic aeroacoustic problems.

TOM JACKSON

Algebraic instabilities in Blasius boundary layer

Algebraic instabilities arise when some initial disturbances, owing their presence to a finite level of noise present in any flow, grow sufficiently to trigger nonlinear mechanisms or to provide new basic states for secondary instabilities. These instabilities are distinguished from exponential instabilities, where infinitesimal disturbances always grow exponentially in time. The presence of algebraic instabilities may lead to the so-called "bypass mechanisms". Work is continuing on the evolution of disturbances in the Blasius boundary layer flow. This work offers a means whereby completely arbitrary initial input can be specified and the resulting temporal behavior, including both early time transients and the long time asymptotics, can be determined.

The approach taken to explore the transient dynamics include linear theory and DNS calculations. The bases for the linear analysis are: (a) linearization of the governing equations; (b) Fourier decomposition in the spanwise and streamwise directions of the flow and; (c) numerical integration of the resulting partial differential equations. DNS spatial calculations are performed to compare with linear theory and to explore the nonlinear dynamics. Because the DNS calculations are expensive when trying to compare to linear theory, we are currently writing a 2D version using high order compact schemes to bridge the gap between the fully nonlinear DNS calculations and the calculations of the linear theory. The results provide explicitly both the early time transients and the long time asymptotic behavior of any perturbation. With this knowledge it is then possible to devise means for flow control and it is possible to either delay or enhance disturbances as the need may be. In addition, the important problem of receptivity can also be analyzed within this framework. All linear results are compared to the equivalent spatial problem using DNS.

In the future, we plan to investigate the receptivity of the boundary layer, as well as various possible control mechanisms for the boundary layer.

This work is conducted in collaboration with Ronald Joslin of NASA Langley Research Center, William Criminale of the University of Washington, and D. Glenn Lasseigne of Old Dominion University.

Numerical simulation of mixing enhancement in a hot supersonic jet

Two NASA programs, High Speed Research (HSR) and Advanced Subsonics Technology (AST) have identified engine noise reduction as an enabling technology. A popular approach to noise

reduction is the introduction of streamwise axial vorticity. Although it has been known for a number of years that tabs can have a substantial effect, many of the details and the physical mechanisms involved are unclear. The objective of this study was to assess the effect of varying the number and placement of the tabs and to elucidate the physical mechanism responsible for the increased spreading rate and consequent sound reduction.

A set of numerical calculations was carried out using the compressible three dimensional Navier-Stokes equations with the tabs modeled by pairs of counter rotating vortices. Both "necklace" and "trailing" type vortices were simulated by changing the sense of rotation of the model vortices. Calculations of the flow with the model vortices present showed that the tabs increased the thickness of the jet about 25% compared to the flow without the tabs, depending on the assumed strength of the vortex circulation. A mixing parameter was increased by a factor of about 2.5 by using six tabs. The effect of the sidewall boundaries on the mixing was also determined by repeating the calculation with six tabs with periodic boundary conditions in the cross stream direction. It was found that, with periodic boundary conditions, the mixing increased somewhat. The results of the calculations also elucidated the basic physical mechanism of the interaction of the vortices generated by tabs with a hot jet and that of the increased jet thickening and increased mixing.

Simulations of the interactions of streamwise vortices with jets of square and circular cross section and jets exiting from lobed nozzles are in progress. We are considering adding a turbulence model to the code.

This research was conducted in collaboration with Chet Grosch of Old Dominion University.

FRANK KOZUSKO

Mobius airfoil

An airfoil design based on a Mobius strip surface appears to have good lift and drag characteristic as well as some strength to weight advantages. This research is to design a Mobius based airfoil and compare lift-drag with equivalently sized airfoils.

This a new project since summer 1996. A basic design is just being completed.

This surface posses problems in selection of grid. Learning and trying GRIDGEN grid generating code is the next step. This will be followed by evaluating the airfoil on one or more 3D flow solvers such as CFL3D, OVERFLOW.

This research was conducted in collaboration with Mikhail Gilinsky of Hampton University.

JOE L. MANTHEY

Numerical methods for computational aeroacoustics

Numerical schemes for computational aeroacoustics are studied for application to duct acoustics. Special issues are the implementation of the boundary conditions at the duct walls and out-flow boundary conditions. Many existing higher order finite difference schemes are not time stable and hence are unsuitable for long time integration.

Eigenvalue stability analysis has been performed for high-order explicit as well as compact implicit schemes with the physical boundary conditions applied. It is found numerical damping or filtering is necessary for stability of explicit schemes. In duct acoustics the lined-wall impedance

condition is often given in the frequency domain. Current work is focused on the time domain implementation of the lined-wall impedance condition.

Future work is to further investigate the numerical stability of time domain implementations of lined-wall impedance conditions.

JAMES E. MARTIN

Flap side-edge noise: Acoustic analysis of Sen's model

Sound generated at the side edges of flaps is a very important, in some cases the most intense source of airframe noise. Recently a new physical mechanism for the flap edge noise source was suggested by Rahul Sen of the Boeing Company. Sen's model suggests that the trailing vortex, existing just off the edge of the flap, may undergo periodic oscillations in the presence of the flap, making the vortex a potentially persistent source of sound generation. This model has several intriguing mathematical aspects and also appears to have physical plausibility based upon flow visualization tests carried out at NASA Langley Research Center. In this study, we numerically investigate the noise generating potential of Sen's model of the flap edge flow field. We investigate the vortex trajectories associated with the model, relate the model parameters to typical aircraft parameters, and propose noise reduction possibilities.

In the model, we consider the flap to be a rectangular slab in the presence of which there exists a potential flow as well as a vortex representing the trailing side edge vortex. Using the Kirchhoff-Routh path function, we have obtained an expression which provides all possible vortex trajectories for a given value of the governing flow parameter. It is only for a limited range of this parameter that the closed vortex orbits of interest occur. Knowing the time dependent vortex motion, the hydrodynamic pressure on the flap surface follows from Bernoulli's theorem. Sound radiation may then be calculated in several different ways. One possible approach is to use a method of matched asymptotic expansions as in earlier studies done on mixing layer noise (Hu, Martin, and Hussaini, 1996). In this study, we instead integrate the Ffowcs Williams-Hawkings equation to obtain the acoustic pressure in the far field. We find the noise generation to depend upon a single characteristic parameter which we relate to typical aircraft parameters. Periodic dipole sound containing the fundamental frequency and several non-negligible harmonics are produced. Our results indicate that the intensity of this sound can be reduced by reducing the chord length or increasing the thickness of the flap.

We are presently studying an analogous three-dimensional model in which the flap edge vortex is represented by a three-dimensional vortex filament.

This research was conducted in collaboration with Jay C. Hardin of NASA Langley Research Center.

D.T. PAPAGEORGIOU

Instabilities in industrial processing flows

Many industrial processes make use of various flow regimes in order to achieve efficient results in such operations as coating and spraying with applications in injection systems or processing of materials. A fundamental understanding of the mechanisms involved is of importance in the

design of more efficient processes or in process control. In this work we have been working with two prototype problems: (i) convective instabilities in microwave heated fluids, and (ii) instability of bi-component jets. The latter problem has applications in processes where jet instabilities may need to be suppressed or enhanced with the resulting product being the formation of two phase spherical particles of different sizes. Applications include printing, pharmaceutical design of slow release pills and crop spraying.

The approach is a combination of modeling, linear and nonlinear stability analysis and numerical simulations. We have developed solutions that predict how microwave heated layers of fluids become unstable under different conditions. One interesting aspect of the analysis is that it predicts a way of cell-size selection as the microwave frequency is varied. A useful feature of microwave induced instabilities in fluid layers is the achievement of similar stability characteristics at different frequencies and powers. We have also studied the absolute/convective nature of co-axial jets and have delineated regions in parameter space (usually the Weber number is used which is a measure of the jet velocity) where the flow changes from absolute instability (relatively low Weber number) to regions of convective or spatial instability at higher Weber numbers. In addition we have developed nonlinear models which are capable of predicting the pinching phenomenon seen in single phase liquid jets.

Future plans include the incorporation of more complicated fluid systems (with temperature dependent conductivities) and a bifurcation analysis to study the nonlinear response. In addition we plan to analyze the singularity formation of the nonlinear models of co-axial jets using the methods developed by us earlier (see the Report in 1995). We are also planning to perform co-axial jet experiments to compare with theoretical predictions of drop sizes and breakup times.

J.R. RISTORCELLI

Compressible turbulence modeling

Primary concern is with the creation of a consistent set of models for the effects of compressibility. The turbulence in supersonic and apparently, hypersonic flows has a low turbulent Mach number. A systematic perturbation procedure applicable to low M_t^2 flows in "compact" flow domains is used.

Focus has been on the reversible transfer due to the 1) pressure-dilatation. The theoretical aspects of the effects of the pressure-dilatation are complete. An algebraic stress closure has been derived using of Girimaji's extension of the ASM theory. The results of the analysis have been tested in a mixing layer and a sizeable decrease in mixing layer growth is predicted. Additional (and larger) decreases in the mixing layer growth are expected to come from the reduction in the 2) shear stress anisotropy. The reduced shear stress anisotropy, due to compressibility, appears to be the most important mechanism in flows compressible shears with nominal transport effects. This effect is associated with the pressure strain correlation. Theoretical progress has been made at the second-order closure level to obtain the effects of compressibility on the shear anisotropy. Further results will await the completion of an exceedingly complex mathematical derivation. A closure for the compressible strain is expected. 3) Most compressible DNS have been conducted using an arbitrary set of initial conditions. In many cases this leads to the superposition of an arbitrary wave field on the DNS which causes excessive levels of dilatational dissipation – substantially larger than those generated by turbulence. The DNS with the initial conditions provided by the theory

is consistent with the theory and with the engineering of compressible turbulence. No effort was expended on the compressible dissipation due to its apparent lack of importance in high R_t and low M_t flows with turbulence production mechanisms as predicted by the theory and substantiated by the DNS.

Future plans in this area include the investigation of the practical and computational implications of the analytical work. These issues are being explored in collaboration with S. Thangam. A substantial portion of the work will be theoretical – at the second-order closure level to account for reduction of the shear anisotropy by compressibility.

Some of the theoretical developments in compressible turbulence are being made possible by collaboration with G. Blaisdell. Interactions, implementations and calculations are being done with S. Thangam, P. Pao, J. Carlson, and H. Khaled.

Aeroacoustics

A pseudo sound analysis was used to obtain the variance of the dilatational field from $k - \epsilon$ and second-order closures. This allows prediction of the near field sound source intensity from RANS type turbulence models. Profiles of the variance of the dilatation associated with the broad banded portions of the turbulence field have been calculated in the over aerofoils. This work continues in conjunction with RANS calculations conducted by M. Santrik of HTC in a program for the Aerodynamic and Acoustics Methods Branch (Thomas/Macaraeg). Similar calculations are being done with S. Thangam in jets and mixing layers.

ROBERT RUBINSTEIN

Transport coefficients in weakly compressible turbulence

The generation of sound waves in compressible turbulence is an effect with no analog in incompressible turbulence. The effect of sound waves on turbulence properties like the eddy viscosity is of interest in connection with turbulence modeling.

A theory of transport coefficients in weakly compressible turbulence is derived by applying Yoshizawa's two-scale direct interaction approximation to the compressible equations of motion linearized about a state of incompressible turbulence. The result is a generalization of the eddy viscosity representation of incompressible turbulence. In addition to the usual incompressible eddy viscosity, the calculation generates eddy diffusivities for entropy and pressure, and an effective bulk viscosity acting on the mean flow. The compressible fluctuations also generate an effective turbulent mean pressure and corrections to the speed of sound. Finally, a prediction unique to Yoshizawa's two-scale approximation is that terms containing gradients of incompressible turbulence quantities also appear in the mean flow equations. The form these terms take is described.

Investigation of heat transfer in compressible turbulence by these methods is in progress. The investigation will provide a simple model for the effects of compressibility on the turbulent heat flux.

This research was conducted in collaboration with Gordon Erlebacher of Florida State University.

S.S. SRITHARAN

Implementation of optimal fluid flow control ideas to NASA's missions

The past ten years have seen an explosive development of the mathematical theory of deterministic, stochastic, and robust control of fluid dynamics. This subject has tremendous potential to impact several aspects of NASA's mission such as internal and external aerodynamic optimization, active real time estimation and feedback control of turbulence, as well as mathematically related subjects such as active control and optimal design of antennae. The objective of this work is to implement the theoretical ideas to certain critical areas of interest to NASA. In particular (a) control of Reynolds' average equations for turbulence and active control within the framework of $K - \epsilon$ models, (b) design of numerical schemes for feedback control of separated flows, and (c) achieve control of electromagnetic waves (with possible turbulence interaction) for the purpose of antenna design and control.

Control with $K - \epsilon$ models is being investigated in collaboration with Dr. Y. Zhou and as a first step we have been looking at cases where control is only applied at the mean level. The question of applying control theory prior to or after the Reynolds averaging is investigated. If the control theory is applied prior to Reynolds averaging, then there are a number of terms to be modeled and they differ from the terms to be modeled if the control theory is applied after averaging. However, the issue of feedback seems to dictate this choice. In collaboration with Dr. J. Loncaric and Dr. B. Allen, approximation of the Hamilton-Jacobi ("the master equation for feedback"), in particular the first order feedback equation (the Riccati equation in infinite dimensions) is used in conjunction with the nonlinear Navier-Stokes equation. We have been comparing several numerical methods to solve Riccati equations in extremely large dimensions. One particularly promising algorithm is the Chandrasekhar algorithm used earlier in smaller dimensions by John Casti, Thomas Kailath, and others. We have also been discussing with NASA Scientists such as R. Wlenzier to address the particular aerodynamic control problem of interest to NASA. In collaboration with Dr. E. Arian, a H^∞ -control algorithm is being designed for active control of electromagnetics. We are also investigating successive approximation methods to obtain linear and nonlinear feedback control laws from Hamilton-Jacobi with specific applications to control of fluid dynamics.

Our future goals are to proceed with these three projects computationally to develop software to be used by NASA Langley Scientists. Reynolds averaged equations will allow us to design feedback controllers for turbulent flows, while developments of efficient algorithms for Riccati equations in very large dimensions will be the fundamental step in active robust (H^∞ -control) design of feedback controllers for fluids and electromagnetics.

L. TING

Turbulent boundary layer, acoustic and structural interaction

The interaction of panel oscillation with the boundary layer and incident acoustic waves (simulating jet noise) is essential for the prediction and control of panel fatigue and the transmission of external noise through panels of an airframe into the interior.

We study the implications of experimental data obtained by Dr. Lucio Maestrello of NASA Langley Research Center on the interaction of incident acoustic waves, turbulent boundary layer

and panel oscillations and the effectiveness of an active control device on the panel oscillation and the transmitted wave. It was found that the peak levels of the power spectra of panel response and transmitted wave increase by an order of 20dB when there is an external pure-tone sound and that a single controller acting on the panel can reduce the peak levels back to the order of magnitude of the case without the external sound. To understand the physics of the complex phenomena described by the experimental data, we attempt to formula simple mathematical models which may simulate some aspect of the phenomena and formulate a numerical scheme simulating of the complex interaction problem using appropriate turbulence model for the boundary layer, nonlinear plate equation for the panel response with a concentrated load simulating the controller and a convective wave equation for the acoustic field outside the boundary layer. For the numerical simulation, we decide to adopt two recent numerical programs developed by Dr. A. Frendi, one for acoustic structural interaction without flow and the second one for the interaction of flexible panel and turbulent boundary layer in a supersonic stream. We need to combine these two programs and modify the second one to handle low subsonic flow.

We intend to construct weakly nonlinear solutions of our simple model equations to explain different aspect of the complex phenomena and to develop a numerical program and then carry out numerical simulation of the phenomena. We intend to continue our earlier theoretical study published recently, "Acoustic field in unsteady moving media", J. Acoust. Soc. Am. **99**, March 1996, by developing numerical programs to implement the rules and formulas presented in this paper to compute the acoustic field induced by a time-dependent source distribution in an unsteady moving medium, whether subsonic, transonic, or supersonic.

This research has been conducted in close collaboration with Dr. Lucio Maestrello of NASA Langley Research Center. To develop the numerical program for the simulation of the boundary layer/acoustic/structure interaction and control, active assistance of Dr. A. Frendi of Analytical Services and Materials, Inc., Hampton, VA is needed. To develop numerical programs and carry out examples implementing the rules and formulas in our paper, active collaboration of Dr. F. Bauer of Courant Inst. Math. Sc., NYU is needed.

GEORGE VAHALA

Lattice Boltzmann approach to fluid turbulence

Lattice Boltzmann techniques (LBE) are a recent attempt to solve macroscopic (nonlinear) problems using mesoscopic methods which are ideally suited for parallel computer architecture. Problems that have been tackled using LBE methods range from Navier-Stokes turbulence, chemical reaction-diffusion systems, multi-phase flows as well as polymeric dynamics. The traditional approach to solving these problems has been through some finite differencing or finite element methods on the actual macroscopic equations. However, the speed of these codes rapidly decreases as the flow geometry complexity increases. Full molecular dynamic (MD) codes, on the other hand, are very computationally intensive and can successfully probe the phenomena only for relatively short length and time scales. LBE is a mesoscopic method straddling MD and macroscopic physics. LBE solves a linearized kinetic equation on a discrete lattice, so chosen that on taking a Chapman-Enskog expansion and the appropriate moments one recovers the macroscopic (nonlinear) equations of interest. However, instead of following the full complexity of the molecular dynamics, one restricts the collisional operator to that which is needed in order to recover the (given)

macroscopic physics. In particular, the standard collisional operator employed in LBE is just the simple BGK linear relaxation operator. What is gained is the simplicity of enforcing complicated geometric boundary conditions without any loss in computational efficiency. This is accomplished by considering the evolution of the LBE distribution function from a molecular dynamic viewpoint. Moreover, it is straightforward to move from 2D to 3D problems. These two properties are sorely lacking in the conventional finite differencing and finite elements methods. It is the objective of this work to explore the limitations of LBE techniques – especially as they apply to thermo-fluid problems.

Our LBE code is 2nd order accurate spatially and 1st order accurate in time. LBE has been tested extensively with the more conventional spectral codes with excellent results for incompressible flows. We have now extended LBE theory to correctly recover the density-independent transport coefficients. Our hexagonal thermo-LBE code has been used to solve the effect of 2D double shear turbulence on a sharp temperature front and to compute the eddy transport viscosity and conductivity coefficients. The code, on a dedicated C90 (at NERSC) has been timed at a concurrency of 15.6 CPU in the 16 processor machine, with over 580 MFlops/processor and a vector length of 127.67.

In the future, we plan to explore algorithms that yield unique equilibrium distribution functions and explore techniques which decouple the spatial and velocity grids. This will enable LBE to handle non-uniform spatial grids as well as high Reynolds number flows.

This research was conducted in collaboration with L. Vahala of Old Dominion University, P. Pavlo of the Czech Republic, and Min Soe of The College of William and Mary.

ROBERT V. WILSON

Simulation of complex, three-dimensional turbulent jets

Three-dimensional, turbulent jets issuing from elliptic or rectangular nozzles exhibit many complex phenomena including strong azimuthal instabilities, switching of major and minor axes, and increased entrainment rates leading to increased mixing. The objective of the present work is to perform numerical simulations of these flows in order to understand such phenomena.

A numerical formulation has been developed which uses higher-order, compact finite differences for all spatial derivatives and a low-storage, Runge-Kutta formulation for time advancement. We have successfully performed spatially-developing simulations of incompressible rectangular and elliptic jets. Direct numerical simulations (DNS) are performed at low Reynolds number and large eddy simulations (LES) at higher Reynolds number. Discrete mode forcing from linear theory or a random, broad collection of modes is used to promote unsteadiness and to study the effect of inlet conditions. Results show that many experimentally observed phenomenon are captured in the simulations such as shear layer roll up, formation and pairing of vortex rings, and axis-switching.

Future efforts will focus on computing statistical quantities and comparing the results to experiment. Large eddy simulations with random forcing functions at the computational inflow will also be performed.

This research was conducted in collaboration with A.O. Demuren of Old Dominion University and M.H. Carpenter of NASA Langley Research Center.

S.L. WOODRUFF

Large-eddy simulations of non-equilibrium turbulent Kolmogorov flow

Large-eddy simulations have been successful in reproducing a variety of turbulent flows, primarily equilibrium flows whose statistics are constant in time. We are concerned here with applying LES to non-equilibrium turbulent flows and particularly with the applicability of various sub-grid models to non-equilibrium turbulent flows. As a test problem we choose Kolmogorov flow, where the Navier-Stokes equations in an infinite domain are driven by a spatially-periodic, deterministic, body force. An instantaneous change in the amplitude of this force sets up a non-equilibrium flow. In this research, the faithfulness with which several subgrid models reproduced this flow were evaluated by comparison with direct-numerical simulation results.

The large-eddy and direct-numerical simulations were performed using a spectral code which employs Fourier modes in the three spatial dimensions and a predictor-corrector method for time stepping. For the large-eddy simulations, the sub-grid Reynolds stresses were added as predicted by the Smagorinsky and dynamic models, as well as the history-integral model introduced by Woodruff. The temporal evolution of a number of global quantities were examined under forcing at several different wavenumbers and the performance of the models was evaluated. The history-integral model did significantly worse than the dynamic model, which itself reproduced the DNS results only for an initial time period. The Smagorinsky model with a constant consistent with the history-integral model gave results similar to the history-integral model; the Smagorinsky model with a constant consistent with the dynamic model gave results similar to the dynamic model. These results indicate that history effects do not play a large role in this problem nor do effects captured by the spatially and temporally varying model coefficient of the dynamic model. Overall, the LES results became poorer as the wavenumber of the forcing increased.

The results of the present study indicate a need for an improved sub-grid model for non-equilibrium flows, at least for high shear and long times; such models are currently being sought.

This research was conducted in collaboration with J.V. Shebalin of NASA Langley Research Center and M.Y. Hussaini of Florida State University.

BASSAM YOUNIS

Development of a non-linear model for the turbulent scalar fluxes

The aim here is to improve the way in which the turbulent scalar fluxes are modeled, particularly in complex strain fields where standard gradient-transport hypotheses are known to fail.

The approach, pursued in collaboration with Tim Clark (Los Alamos National Laboratory) and Charles Speziale (Boston University), has been to use representation theorems to expand a functional relation for the turbulent scalar fluxes. The outcome is an expression for the scalar fluxes which is both explicit in these quantities and is directly dependent on the Reynolds stresses and the spatial gradients of mean velocity (as required by the exact transport equations). The model coefficients have been evaluated from various LES and DNS data bases and the complete model successfully benchmarked against standard test cases.

Further work in progress is aimed at extending this model to handle the difficult case of a heated wall jet developing over a convex-curved surface for application to wing de-icing. This work was conducted in collaboration with John Shebalin of the Aerodynamics and Acoustic Methods Branch.

YE ZHOU

Time correlations and the frequency spectrum of sound radiated by turbulent flows

Theories of sound radiation by turbulence have generally been based on a highly simplified treatment of the time correlations in turbulent flow. It is important to re-evaluate some of the previous predictions for sound radiation by turbulence to assess the effect of these simplifications.

Theories of turbulent time correlations are applied to compute frequency spectra of sound radiated by isotropic turbulence and by turbulent shear flows. The hypothesis that Eulerian time correlations are dominated by sweeping by the most energetic scales implies scalings for the sound radiated by isotropic turbulence which are different from predictions of a classical theory of jet noise based on dimensional analysis. It is shown that the classical scalings are obtained by simplifying the description of turbulent time correlations. An approximate theory of the effect of shear on turbulent time correlations is developed and applied to the frequency spectrum of sound radiated by shear turbulence. The predicted steepening of the shear dominated spectrum appears to be consistent with jet noise measurements.

The model of time correlations developed here is being used to synthesize a stochastic model of the sound source terms for calculation of far-field sound radiation using the linearized Euler equations. These methods are also being applied to compute sound radiation by heated jets.

This research was conducted in collaboration with Robert Rubinstein of ICASE.

Energy transfer in compressible turbulence

This study investigates the compressible energy transfer process. We extend a methodology developed originally for incompressible turbulence and use databases of a numerical simulation of a weak compressible turbulence based on EDQNM closure. In order to analyze the effects of the compressibility, the velocity is decomposed into a solenoidal and compressible part. Using EDQNM modeling, it is then possible to obtain energy transfer equations in the case of a weakly compressible turbulence. In this study, we are interested in the most fundamental building block of the energy transfer process, that is to say the triadic energy transfer. We found that the solenoidal triadic energy transfer is not affected by the compressible effects. The study of the compressible triadic energy showed that the compressible mode receives energy from the solenoidal part and that the energy is inputted and removed locally.

A more accurate study of the compressible triadic energy transfer has been done: a decomposition and a study of all the solenoidal and compressible triadic transfer terms permitted us to demonstrate that an energetical cascade appears on the compressible mode for high turbulent Mach numbers. Furthermore, a publication for the Journal of Fluid Mechanics has been finished.

In the future, it would be interesting to study the scaling of triadic interactions and the energy flux to derive a subgrid scale model.

This work has been done in collaboration with Françoise Bataille of the Institut National des Sciences Appliquées de Lyon, France.

Subgrid-scale models for turbulent flows

In large eddy simulations (LES) several different competing models are currently in use. While these models successfully capture some features of turbulent flows it is well known that they are deficient on several accounts. For instance, correct predictions of energy backscatter from the subgrid

scales to the resolved ones present a significant challenge for the existing models. Deficiencies of the models may be traced to our inadequate knowledge of nonlinear interactions between resolved and subgrid scales requiring various ad hoc assumptions about these interactions. Better models can be expected only if more accurate information about nonlinear interactions in turbulent flows is used. Such information is available from analyses of databases obtained in direct numerical simulations of turbulent flows and is currently being employed in the design of improved subgrid scale (SGS) modeling procedures. The objective of this research is to implement the new SGS model for a specific flow and compare its predictions with the exact SGS quantities in *a priori* tests.

The basic concept in the proposed SGS modeling procedure is to generate a range of subgrid scales from the resolved scales using properties of the LES filtering operation and the nonlinear interactions. This procedure has been implemented for isotropic fields in periodic domains. Using velocity fields generated on a 128^3 mesh by an incompressible, pseudo-spectral Navier-Stokes solver several SGS quantities were calculated exactly without any model and with the use of the model. The computed quantities were SGS stress tensor, SGS force, and SGS energy transfer. To assess the performance of the model the statistical moments (mean and rms values) of the exact and modeled quantities as well as the correlation coefficients between them were calculated. The modeled mean and rms values are generally within 5% of the exact values and the correlation coefficients exceed 70%. This agreement is much better than observed for classical SGS models where the correlations usually do not exceed 30%. This implies that the new model has a potential to significantly improve predictions of turbulent flows.

In the near future we plan to perform similar analyses for several more velocity fields and filtering procedures as well as implement the model in time evolving large eddy simulations.

This research was performed in collaboration with J.A. Domaradzki of the University of Southern California and P.K. Yeung of Georgia Tech.

Strongly stratified limit of Boussinesq equations

The turbulent flows that are subject to stratification have many important applications in geophysics and engineering. Experimental studies, direct numerical simulations, large-eddy simulations, and closure approximation have established that stratification suppresses the nonlinear energy cascade from large scale to small scale. The effect of stratification is through phase scrambling. This evidence formed the basis for extending the phenomenology originally developed for rotating turbulence to the stratified turbulence. The objective of this work is to achieve an understanding of the nonlinear physics of gravity mode/vortex mode interaction to serve as a guide in the development of better turbulence models.

We focus our attention on developing a better understanding of the strongly stratified limit of Boussinesq equations. We have deduced the energy spectrum and spectral eddy viscosity that have an explicit dependence on stratification rate. Our analysis is based on the fact that this is a two time scale problem. We first present a general mathematical framework based on the Green's method which is applicable when there are disparate time scales. The approach leads to the Poincaré transformation and the Poincaré variables. The transformed equations are used to obtain the time-dependent dynamical theorem valid in the asymptotic limit of strong stratification. It shows that the dynamics of vortical and wave components of total field decouple in the strongly stratified limit. This fact allows to obtain mathematically rigorous equations describing vertical

variability which are exact in the asymptotic limit of strong stratification.

In the future, we plan to investigate the combined effect of rotation and stratification on homogeneous turbulence.

This research was conducted in collaboration with A. Mahalov and B. Nicolaenko of Arizona State University.

Mixing and transport in rotating turbulence

Transport processes in rotating turbulent flows occur in important applications such as the mixing of chemical species in jet engines and pollutant dispersion over large spatial dimensions in the atmosphere. However, the underlying physics is little known, and current models are (as in the case of mean flow predictions) incapable of accounting for the effects of rotation in a non-inertial frame of reference. The objective of this work is to provide fundamental understanding in the processes of passive scalar transport and fluid particle dispersion in rotating turbulence, as a basis for the guidance of modeling.

Direct numerical simulations at 64^3 resolution have been carried out for initially isotropic turbulence subjected to uniform solid-body rotation, using a pseudo-spectral parallel computer code on an IBM SP2. Analyses of spectral scalar transfer and Lagrangian statistics have been performed using data archived in the simulations. The results indicate that, similarly to the velocity field, the scalar fields develop a quasi- two-dimensional structure. The mixing of passive scalars is reduced by rotation, which is traced to a weakened spectral cascade towards the small scales and a consequent reduction of the molecular dissipation of scalar fluctuations. Single-particle Lagrangian statistics indicate that material fluid velocity components remain correlated for much longer periods along the axis of rotation. This leads to a departure from classical theories of turbulent diffusion. Acceleration correlations exhibit damped oscillations which reflect the nature of the Coriolis force acting on the fluid. All results obtained (while still preliminary) indicate that rotation has a dramatic effect on important characteristics of mixing and transport, and hence point to the need for the development of new models.

Higher-resolution computations are planned in the near future in order to achieve higher Reynolds numbers (up to 150 based on the Taylor scale, on 256^3 grids). In future work we also plan to quantify the effects of different molecular diffusivities and to study two-particle statistics which describe the dispersion of fluid particle pairs. In view of increased statistical variability due to rotation in many of the results in this work, ensemble averaging over multiple independent realizations will be required.

This work was performed in collaboration with P.K. Yeung of the Georgia Institute of Technology.

Turbulence modeling for strongly-stabilized rotating shear flows

The flow in a pipe rotated about its longitudinal axis has proven difficult to calculate. In eddy-viscosity models, for example, rotation terms are absent from the model equations and the effects of rotation are therefore missed entirely.

Corrections to the production and destruction terms in the dissipation rate equations and modification of the eddy viscosity have been proposed based on theoretical considerations. These corrections were evaluated and refined using data from experiments and large eddy simulations. It

was found that the effects of strong rotation leading to collapse of turbulence and the reversion of the mean axial velocity profile to a nearly parabolic distribution were well predicted by the new model. Combination of the new model with previously proposed nonlinear constitutive relationships yields profiles for the tangential velocity in close agreement with data.

The extension of this type of modeling to other cases of rotation and swirl, including destabilizing effects, is under investigation.

This research was conducted in collaboration with R. Rubinstein of ICASE and B.A. Younis of the City University of London, England.

Unsteady RANS input to computational aeroacoustics

The purpose of this work is to predict the aeroacoustic field associated with shedding of vortices from a circular cylinder in turbulent flow conditions.

The strategy being followed is to use unsteady Reynolds-averaged Navier-Stokes (RANS) calculations with turbulence models developed at City University to calculate the time-dependent distribution of the static pressures on the wall. The computed pressure values are currently being used to provide an input to NASA Langley's aeroacoustics code (activity with Ken Brentner - Aerodynamics and Acoustic Methods Branch).

Further work is in progress with B.A. Younis to discover whether further improvements in the modeling of the effects of periodic mean-flow unsteadiness are possible within the context of eddy-viscosity closures.

APPLIED COMPUTER SCIENCE

DAVID C. BANKS

Visualizing vector fields

Vector fields present a challenge for visualization. If we populate a 3D vector field with streamlines, the resulting set of curves creates a dense volume that obscures its own interior. If the vector field is time-varying, how should its streamlines behave? Their constituent points at time t_0 could individually follow the vector field, but the result will not (in general) be a set of streamlines at time t_1 . But alternatively, if we populate the vector field at time t_1 with a new set of streamlines, the result will not (in general) show any visual coherence in an animation. The aim of the research is to explore reasonable visualizations of such a time-varying vector field.

We have investigated several approaches to the visualization problem. Joint work with Tom Crockett, Will Bene, and Bart Singer on visualizing a transitional boundary layer flow led to Bene's masters project (Old Dominion University). In this work, vortex cores (and the surrounding vortex shapes) are interpolated to animate coherent structures deduced from a direct numerical simulation of a young turbulent spot. Because the cores approximate integral curves of vorticity, it is reasonable to advect them point-wise. Streamlines of velocity are not so well behaved. Work with Greg Turk (UNC and Georgia Tech) led to a technique to evenly place streamlines in 2D vector fields. This work is being continued at Mississippi State to animate such integral curves in an time-varying vector field. Our previous work at ICASE on illuminating curvilinear reflectors has led to a new method for illuminating clusters of streamlines within a flow volume. Finally, we have developed a technique for creating multi-frequency oil-streak images in 2D vector fields.

Because time-varying 3D flows require enormous amounts of storage, we are investigating multi-resolution representations of the features within them. We plan to extend our 2D streamline placement techniques to 3D vector fields. We plan to extend the streamline placement and the computational oil-streak techniques to time-varying flows.

SHAHID H. BOKHARI

Bounded contention complete exchange on meshes

Complete Exchange is an expensive operation on mesh connected machines such as the Intel Paragon. Prior research by David Scott shows that, on an $N \times N$ node machine, this operation takes $\Theta(N^3)$ steps. In actual practice the time to synchronize processors adds substantially to the total time. The objective of this research is to develop fast algorithms for complete exchange that minimize the impact of link contention and synchronization overhead.

We have developed a generalization of Scott's algorithm that allows us to trade off synchronization time against communication overhead. Our approach is to permit a controlled degree of link contention to take place by reducing the number of synchronization operations. This is made possible by a carefully designed *collapsible* schedule of communication steps. An algorithm to generate the collapsible schedule has been developed and tested on the "Trex" 512-node Paragon mesh at Caltech.

The theory behind collapsible schedules for 2- and 3-dimensional meshes needs to be placed on a solid foundation. Depending upon the availability of the machine, we may test this algorithm on the ACSI mesh, which has a dual mesh interconnect.

This research was conducted in collaboration with David M. Nicol of Dartmouth College.

High performance communication hardware for parallel computers

Communication overhead plays an important role in limiting the performance of parallel computers. The powerful "wormhole" communication mechanisms available on today's parallel machines have the potential to allow us to communicate from one processor to many other processors in one step. When used in an intelligent fashion, this permits us considerable performance enhancements.

This research aims at identifying the hardware configurations that are required to obtain performance improvements in the "complete-exchange" communication pattern. For example I can show that when the complete exchange is executed with an intelligent wormhole routing mechanism and a carefully developed schedule, the number of message startups can be halved. This can lead to major savings in communication time.

I plan to study the architecture and applications of high performance communication hardware on 2- and 3-d meshes. Hypercubes will also be studied since prior research has shown that contemporary parallel machines can execute hypercube specific algorithms reasonably well.

GIANFRANCO CIARDO

Kronecker operators for the description and solution of Markov models

The solution of large Markov models is limited by the size of the transition rate matrix R . Even when stored in sparse format, the practical limit for a modern workstation is substantially less than 10^6 states and 10^7 nonzero entries, if we want to avoid the use of virtual memory. Many interesting problems in performance and reliability, however, correspond to much larger models. Our objective is to increase by one order of magnitude these limits.

To achieve this goal, we use Kronecker operators (also known as tensor operators) for the description of R . We decompose the model into K interacting submodels, each described by a (small) transition rate matrix $R^k, k = 0, \dots, K - 1$, and express R as the Kronecker sum of the matrices R^k , plus some Kronecker products of "corrective factor" matrices, also of small size. Unlike previous approaches, we focus on methods that require vectors having only as many entries as the number of actually reachable states in the model. This requires the reachability set to be generated and stored efficiently. To date, we have: (1) described how to manage models that have an extensive type of state-dependent behavior, both in the structural and stochastic aspects, and shown how to manage instantaneous synchronizations between submodels; (2) defined an efficient multilevel data structure for storing and searching the reachability set. Our prototype implementation can store over ten million states on a single workstation, and has a state search time much lower than with traditional storage approaches based on a single binary tree; and (3) derived an algorithm, based on the above data structure, which greatly reduces the complexity overhead inherent with Kronecker operators.

We plan to explore and implement an overall solution approach that integrates Kronecker operators and the efficient storage of the state space in a single distributed algorithm. We envision

being able to solve models with 10^8 states and 10^9 nonzero entries using a dozen processors having 128 megabytes of main memory each.

THOMAS W. CROCKETT

Parallel graphics libraries for runtime visualization on massively parallel architectures

Applications which run on large-scale parallel computing systems often generate massive volumes of output data. To assist in the analysis and understanding of this data, we are developing parallel algorithms and software which allow application programs to generate visual output at runtime.

The focus over the last several months has been on documenting our PGL graphics library and preparing it for release to the HPCC community by the end of 1996. A pre-release version of PGL has been adopted for use in two new applications: experimental evaluation of a multi-ported frame buffer at Princeton University, and visualizing the results of turbulent transport calculations at Caltech. We anticipate that release of a robust, documented version of PGL will enable many additional researchers to incorporate parallel rendering technology into their large-scale parallel computations. The experience we have gained in designing and using PGL provided the basis for an invited paper on software architectures for parallel rendering, which we presented as a keynote address at the First Eurographics Workshop on Parallel Graphics and Visualization in Bristol, England.

In the near term, we are planning ports of PGL to workstation networks and newer parallel architectures such as the Cray T3E. We also plan to incorporate additional features such as transparency, a sphere renderer, and dynamic color quantization. We have several algorithmic avenues which remain to be explored, including the important issue of scaling to large numbers of processors (hundreds to thousands).

Lossless image compression for parallel rendering applications

In recent years, a significant amount of research has been done on scientific visualization on large-scale parallel machines. Comparatively little work has been done on the problem of transmitting the images out of the parallel machine via communication networks onto a user's desktop where they must be viewed at the rate of 0.2 to 10 frames/sec. The parallel lossless compression time, the communication bandwidth and the decompression time together determine the resulting throughput. The objective of this work is to understand the characteristics of the existing lossless compression algorithms and develop new compression algorithms suitable for parallel rendering applications.

We conducted experiments to measure compression effectiveness, compression time and decompression time for several well-known lossless compression algorithms (Huffman, arithmetic, LZW, dmc, ppm, rle and ymf) on image data on Mars and Landsat earth terrain data to identify potential candidate algorithm for parallel implementation. We also proposed new lossless compression algorithms based on block decomposition of the color channels of the images. The significance of this work is that we now have data to identify the parameters that will affect the choice of the algorithm and the architecture for parallel rendering applications in the context of users served by a parallel machines over communication networks.

In the future, we plan to implement the new lossless compression algorithms, select a couple of the most suitable algorithms for parallel implementation and measure the performance of the overall system for NASA applications.

This research was conducted in collaboration with A. Mukherjee of the University of Central Florida.

STEPHEN GUATTERY

Multilevel separator algorithms and applications

Graph separators have important applications in numerical computation for tasks such as partitioning and computing elimination orderings via nested dissection. Since finding optimum separators under useful definitions of "optimum" is NP-complete, separator algorithms are often evaluated on the basis of their performance in a particular application. A recent trend has been the use of multilevel separator algorithms. These algorithms work in a way roughly analogous to multigrid algorithms: a graph is coarsened in several steps to produce a smaller graph, which is partitioned. The graph is then uncoarsened; at each uncoarsening step the lower-level cut is extended to the finer graph, and possibly refined. Recent algorithms by Karypis and Kumar, Ashcraft and Liu, and Hendrickson and Rothberg have followed this general model; there is wide variation in how they work (e.g., computing edge versus computing vertex separators; separators versus multisector approaches). These algorithms are not yet well understood with respect to their performance on particular classes of graphs, nor with respect to how they work in various applications. Gaining such an understanding will allow people to choose the most efficient partitioners for their applications.

We have recently made some steps in this direction. Pothén and Kurfert at Old Dominion University had observed that METIS, which uses a multilevel edge separator algorithm as part of a nested dissection routine, produced an ordering for an example graph (ken13) that required three orders of magnitude more work in Cholesky factorization than did the ordering produced by the MMD heuristic. A similar experimental multilevel nested dissection algorithm produced an ordering that required between two and three times more work than MMD. We examined the ken13 graph and the two algorithms, and determined why their performances were different. The graph in question is from a linear programming problem, and specifically from a multicommodity flow problem. We also determined that METIS found a small edge cut, and that the vertex separator associated with this produced much fill when used in nested dissection. The experimental algorithm, on the other hand, had a similar structure to METIS, but used a vertex separator algorithm at the lowest level, and a refinement method that considered vertex separators. These results showed that one must be careful in making assumptions about the structure of graphs arising in applications, and also showed some aspects of the performance of two multilevel algorithms.

We plan to continue and extend this work. We have done some preliminary work toward an analysis of the performance of multilevel coarsening schemes. Our goal is to determine if there are classes of graphs for which these schemes produce separators of bounded size, and to specify properties that make multilevel algorithms good for particular applications.

This was joint work with Alex Pothén of ICASE and ODU.

MATTHEW HAINES

OpenThreads: An open implementation of lightweight threads

Most lightweight thread packages are designed and implemented using the traditional *black-box* approach which hides the implementation details and provides a simple user interface. However, black-box abstractions do not always work, because there are times when the implementation strategy for a system cannot be determined before knowing how the system will be used in a particular case. This is particularly true for runtime system designers who want to make their own implementation decisions regarding thread states and transitions, thread lists and scheduling algorithms, context switching modes, stack management, timing and profiling, etc.

In contrast to the traditional black-box design, the design of OpenThreads is based on an *open implementation* where critical implementation decisions are exposed to the system-level programmer in a clean and controlled manner. OpenThreads was designed to satisfy three essential goals: *flexibility*, *efficiency*, and *portability*.

Within the past year we have finished the initial design and implementation of OpenThreads, and have tested the software by using it to implement the Orca concurrent programming language under two operating systems: Solaris and Amoeba.

We plan to continue or evaluation of OpenThreads by using it to construct other runtime systems that rely on lightweight threads. We will make modifications to the design as warranted by the evaluations.

This research was conducted in collaboration with Koen Langendoen of Vrije Universiteit, Holland.

VICTORIA INTERRANTE

Visualizing layered surfaces in volume data

Researchers in computational fluid dynamics need efficient and effective methods for visualizing the complicated, three-dimensional time-varying vector and scalar fields that are generated in the course of their numerical simulations. The effective representation of layered structures - such as multiple stream surfaces, shock surfaces, or superimposed isovalue surfaces of pressure or density - in large volumetric datasets presents specific computational challenges, due both to the size of the datasets and to the complexity of extracting and illustrating the essential features in them. Transparency can be a useful device for representing the complex spatial relationships between multiple superimposed surfaces only if the overlapping surfaces are rendered in such a way that the 3D shapes, positions and orientations of both the external and internal surfaces can be clearly seen. The threefold goals of my research are: to define (mathematically) the visually essential elements of surface shape; to identify promising graphical techniques for more effectively communicating the three-dimensional shape of a transparent surface; and to explore methods for more efficiently rendering multiple transparent surfaces in volume data. This work is relevant to a number of different applications here at ICASE.

A 3D rendering of the surfaces of interest in axisymmetric numerical data simulating the interaction of a shock with a longitudinal vortex provides a complementary means of understanding this data and may facilitate appreciation of certain aspects of the spatial geometry. However, a

direct volume rendering of the 3D data will produce an image in which much of the important information is obscured. I have used a two-dimensional ridge detection technique to extract "feature lines" defining the shock and slip surfaces from 2D density images of the flow and have explored methods for comprehensibly rendering these surfaces in 3D.

My previous investigations into perceptually-inspired methods for texturing transparent surfaces with directed short strokes have yielded A novel method for illustrating flow in volume data is obtained by advecting volume particles of both positive and zero opacity along the directions indicated by the flow to produce a 3D grey-scale mask that selectively modulates the opacity with which an isosurface is rendered. Particle tracing is restricted to the relatively small portion of the volume immediately adjacent to the predefined surface, yielding more comprehensible images and making this approach more computationally efficient than full volume line integral convolution. The tracing of "empty" as well as "full" particles minimizes clumping and allows a reasonable distribution of lines to be obtained without any fine tuning. Occluding "halos" can enhance depth cues and variable-length filters and variable-resolution input textures can selectively emphasize important flow information.

There are several directions for future work, both in the human factors arena (to better understand the perceptual capabilities of the human visual system and the impact of various kinds of surface markings on observers' ability to perceive surface shape) and in the implementation of high-performance algorithms for generating perceptually-enhanced images at the nearly interactive frame rates required for practical use.

Various aspects of my work were performed in collaboration with Gordon Erlebacher, Chester Grosch and Kwan-Liu Ma.

JIM E. JONES

Parallel multigrid methods

Standard multigrid methods are not well suited for problems with anisotropic coefficients which can occur, for example, on grids that are stretched in order to resolve a boundary layer. There are several different modifications to the standard algorithm that yield efficient methods for these problems. This project is aimed at developing parallel multigrid algorithms that are robust with respect to anisotropic coefficients.

The research thus far has taken two tracks. First, a semi-coarsening multigrid algorithm for logically rectangular grids has been implemented in High Performance Fortran (HPF) on the NASA SP2. Second, we have implemented a multigrid algorithm based on block-structured grids and line or plane relaxation within the blocks. These block structured grids are commonly used in NASA fluid dynamics codes to deal with complex geometries and facilitate parallel processing. We began looking at a 2D model problem – the anisotropic diffusion equation. Using serial numerical experiments and rigorous analysis, we have derived a relationship between the block size, the amount of overlap between blocks, and the strength of the anisotropy that must hold for the resulting multigrid algorithm to be efficient.

We intend to continue testing of the HPF multigrid code on the SP2. As the HPF compilers mature, we should be able to answer questions about the efficiency of the algorithm and the suitability of the language for multigrid applications. We also intend to continue work on the block-structured

algorithms. In particular, our analysis indicates that the result relating block size, overlap, and anisotropy strength holds in more complex situations – more complicated PDE's and/or more complicated problem geometries. We would like to verify this by numerical experiments as it could provide guidance to practitioners using these block-structured grids.

The work on block-structured grids was conducted in collaboration with Duane Melson of NASA Langley Research Center.

HARRY F. JORDAN

High speed optical sorter

Optical signals have been used extensively in communication, and bit rates up to 100 Gb/s are achievable. It is not clear whether any substantive processing can be done on signals at this rate. Sorting provides a problem that can potentially exploit such high bit rates. The objective of this work is to demonstrate a system design that can do sorting on multiple channels, each operating at 100Gb/s, and demonstrate that useful processing can be done with optical signals at high bit rates.

The key to practical high speed sorting is to separate bit level switching from switching of full words. A radix sort allows control to be derived from a single bit per word. This bit must be sensed by a costly, high-speed device, but the rest of the logic for reordering is done at the word rate. Use of partitioning in the time domain and a modest number of parallel space channels allows single radix bit sorting to be done using a small banyan network and a time-domain compaction structure (packer) in each separate channel. The packer is a pipelined network with distributed control optimized to minimize optical switches and logic. The logic of the optical sorter has been demonstrated using lithium niobate switches running at 50 Mb/s. The high-speed portions of the system have been shown to be feasible at the 100 Gb/s target rate.

We plan to experimentally investigate the high-speed device (TOAD gate) that optically extracts one radix bit from a 100 Gb/s signal and uses it to control subsequent logic. Bit extraction has been demonstrated at this rate, but not integration into a system.

This work is done in collaboration with Kyungsook Lee and Rajgopal Kannan at the University of Denver, Coke Reed at the Center for Computing Sciences, and Jon Sauer and Douglas Straub of the University of Colorado.

Architecture independent metric for the data movement demands of a parallel program

Recent work has shown that an architecture independent metric can be applied to shared memory SPMD style parallel programs to predict their demands for data movement among processors. The metric characterizes the demands at different time scales, and when applied at the time scale of a particular interconnection network latency, it accurately predicts program communication behavior across a wide range of network architectures. This behavior was verified over the SPLASH benchmark set and three different multiprocessor interconnection network architectures. We now extend this work from shared memory SPMD programs synchronized by barriers and critical sections to distributed memory multiprocessors synchronized by send and receive operations.

The key problem in establishing an architecture independent program metric is to establish an abstract time scale across the processors. This is harder with send/receive synchronization than

with barrier synchronization. Local ordering of sends and receives is known to have an impact on program timing, and much time is spent in optimizing communication in such programs. The use of interprocess dependence graphs generated from communication operations offers a mechanism for establishing the abstract time scale. Given the abstract time scale, the amounts of data transferred in each communication will allow a temporal characterization of network demands placed by a given program. The applications of such a metric include architecture independent program restructuring to reduce communication demands and determination of the capacity required of the communication network of a machine to run this program efficiently.

In the future, we will develop the abstract time scale implied by send and receive pairs and use it to define the data movement metric for distributed memory multiprocessor programs. We will test the metric on a range of benchmark programs and compare with timings on different interconnection network architectures.

This work is done in collaboration with Gita Alaghband of the University of Colorado at Denver and Bernardo Rodriguez of the University of Colorado at Boulder.

DAVID E. KEYES

Parallel algorithms of Newton-Krylov-Schwarz type

The crux of Newton-Krylov-Schwarz (NKS) algorithms for the parallel solution of PDEs is a balance of implicit convergence rate and good data locality, which we characterize by the slogan "Think Globally, Act Locally." Thinking globally refers to judicious use of a coarse grid and acting locally refers to accommodation to the pronounced memory hierarchies of modern high-performance machines, in which both vertical (memory/cache) and horizontal (processor/processor) transfer of data exacts a performance price. Our objective is to design and tune parallel domain decomposition algorithms with inputs from both analysis and architecture.

Our approach during the current period has combined theory and practice. Practice led theory in the use of adaptive pseudo-transient continuation for nonlinear problems to which Newton's method cannot be directly applied. In Euler flow problems on unstructured grids, we have identified three stages of pseudo-transient continuation and built a convergence theory upon them. Pseudo-transient continuation is widely used, but sufficient conditions for convergence seem not to have been spelled out before. Theory is leading practice in the development of two-level preconditioners for purely hyperbolic problems. Results previously obtained for optimal convergence for elliptic and parabolic problems ("optimal" referring to the absence of degradation of convergence rate as the mesh is refined or the number of subdomains increases) have been extended to the hyperbolic case. Time step selection plays an important role in the theory. Meanwhile, we continue to implement NKS methods for a variety of applications in CFD, turning attention to software issues (see the following topic).

We plan to perform numerical tests on linear hyperbolic problems to evaluate the new Schwarz theory and to continue to apply NKS techniques to computational aerodynamics applications.

Important collaborators in various aspects of this work have been: Xiao-Chuan Cai of the University of Colorado-Boulder, Tim Kelley of North Carolina State University, and Yunhai Wu of Old Dominion University.

PDE-oriented problem solving environments

The Portable Extensible Toolkit for Scientific Computing (PETSc) being developed at Argonne National Laboratory is a large and versatile package integrating distributed vectors and matrices in several local sparse storage formats, Krylov subspace methods, preconditioners (including Additive Schwarz), and nonlinear solvers (including Newton-based methods and pseudo-transient continuation). Each software component consists of an abstract interface and one or more implementations using particular data structures. This design leads to layered procedures for the various phases of solving PDEs, with a common interface style for each class of problems. PETSc's rapid algorithmic prototyping environment can be converted into a high-performance environment by changing compilation flags, as opposed to source code. Our objective is to use PETSc to parallelize flow solvers at ICASE and LaRC.

Our approach is to work with the PETSc developers to suggest and test new features that serve the needs of computational fluid dynamicists (among others). Our experiences with full potential, Euler, and combustion codes are responsible for many existing and evolving features of the package, which currently runs on more than ten architectures, including ICASE machines and the NASA CAS HPCCP machines, apart from (so far) the Power Challenge array at Ames. A three-dimensional structured-grid implicit Euler code runs in parallel at over 1.2 GigaFlop/s on 32 nodes of an IBM SP2, using Schwarz-preconditioned GMRES.

Our short-term goal is to make the porting process "routine" for codes based on structured grids. In the long term, we will build on this experience in the unstructured grid setting.

Lois Curfman McInnes of the PETSc development team, Dinesh Kaushik and Nilan Karunaratne of Old Dominion University, and Abdelkader Baggag of the University of Minnesota have participated in this work.

KWAN-LIU MA

Extracting features from 3D unstructured meshes for interactive visualization

The three-dimensional unstructured grids used for the numerical simulation of 3D flow are generally very large in size and irregular in both shape and resolution. Even the simplest renderings of many of these meshes can be time-consuming to compute on an average desktop workstation, and once an image is produced it can be difficult to adequately perceive relevant geometric structure through the tangle of overlapping lines.

This research investigates techniques, based on the extraction of a small set of perceptually significant geometric features from the surface mesh, for facilitating the visualization and interactive manipulation of the typically very large and dense three-dimensional unstructured grids used in aerodynamics calculations. Displaying the feature lines in place of the full model may both considerably decrease the rendering latency and at the same time improve the comprehensibility of the presented data. When rendering time is not of critical concern, it can be useful to display feature lines in conjunction with surface or volume-rendered data to highlight essential structural detail of the underlying geometry while preserving the visual prominence of the flow information.

We have developed several simple but efficient and effective techniques for extracting perceptually-relevant feature lines, such as ridge- and valley-like lines as well as silhouette edges, from unstructured triangular meshes. These methods are very fast and require little or no preprocessing of the

data. They allow the clarity of the display to be enhanced while enabling complex models to be interactively manipulated in a fraction of the time that is ordinarily required for a rendering of the complete dataset. In next page, the left image shows the original mesh and the right one shows the feature lines extracted using our methods.

Some of the techniques we have developed require view-dependent calculations. It would be preferable to come up with a fully viewpoint independent technique that is capable of elegantly portraying all essential features at unambiguously interactive rates. We also hope that such algorithms can be incorporated into future commercial systems for 3D CFD data display, as we find them very useful and believe that others will too.

This work is being done in collaboration with Victoria Interrante of ICASE.

Interactive exploration of large 3D unstructured-grid data

In aerodynamics calculations, unstructured grids are used to model objects with complex geometry. Because the grids are typically large in size and irregular in both shape and resolution, often special data processing and rendering algorithms are needed to just make possible visualization of the simulation results. This research studies the needed software support for conducting the desirable iterative, interactive visualization process to analyze very large unstructured-grid data on an average workstation.

We propose a visualization process taking the following two steps. The first step attempts to derive desirable viewing and rendering parameters and to locate regions of interest. This may be performed on a workstation using a fast but less accurate rendering algorithm on a coarse representation, thus a much smaller version, of the original data. Consequently, we need a multi-resolution representation of the original data (mrrd). The mrrd allows interactive exploration of the data. Once a hot spot is identified at a particular resolution, the user may switch to viewing at a higher resolution. This exploration process continues until the region of interest and viewing as well as rendering parameters are completely determined.

The second step takes the parameters derived, extracts the selected sub-volume out of the original data, and invokes a more accurate rendering program to produce high quality visualization results. Note that because of its size, the original data set must be stored on disk. It is thus essential to have adequate database support such that a sub-volume can be quickly retrieved from disk. We represent the data using R-tree.

The mrrd that we have developed in conjunction with a fast splatting rendering method make possible interactive visualization on a workstation. The following four images show rendering of an ONERA M6 wing data set at different resolutions. We have also conducted both analytical and experimental study (using a 4-million tetrahedra data set) comparing R-tree with other data structures; and our results show that R-tree is a better for fast retrieval of disk-resident, region data like tetrahedra. Together, the proposed visualization process allows computational scientists to study their data at the highest possible resolution in a more efficient manner, rather than reducing the data or operating at a very inefficient batch-mode.

We intend to conduct more extensive comparison tests of the data retrieval methods involved. First we plan to implement well balanced octrees for comparison. Next we intend to test the techniques on much larger data sets, 20-500 million tetrahedra.

This work is being done in collaboration with Scott Leutenegger of Denver University and Dimitri Mavriplis of ICASE.

PIYUSH MEHROTRA

Software environment for multimodule parallel applications

Exploiting the multiple levels of parallelism in multimodule applications, such as multidisciplinary design of aircrafts, requires software support beyond what is provided by data parallel languages such as High Performance Fortran. The objective of this project is to build a software system, including language extensions, compilers and runtime system, to support coarse grained parallel tasks which may be internally data parallel across a distributed heterogeneous network of workstations and massively parallel machines.

Based on our experiences with expressing codes in an earlier version of the language, we have simplified the design of Opus while giving the user more control. We have unified the various constructs of Opus into a single one. The central concept is that of a ShareD Abstraction (SDA), a set of data structures along with methods to access the data. Methods can be activated asynchronously or synchronously and each method, when executing, has exclusive access to the data. Thus SDAs can be used both as computation servers for encapsulating individual discipline codes. The methods can be invoked independently providing a framework for asynchronous execution of tasks. Encapsulation allows discipline codes to be changed in a plug-compatible manner as long as it keeps the same interface. The SDAs can also be used as data repositories for communicating and sharing data between the computation servers. This allows any interdisciplinary interpolation of data to be handled by the intermediary SDA without either discipline being aware or responsible for massaging the data.

In the coming months, we plan to implement a web-based user interface for Opus which will allow multiple users to control and monitor the execution of an Opus program across a heterogeneous environment.

This work is being done in collaboration with Matthew Haines of the University of Wyoming, John Van Rosendale of ICASE, Barbara Chapman of the European Center for Parallel Computation, and Hans Zima of the University of Vienna.

Multithreaded system for distributed environments

Traditionally, lightweight threads are supported only within the single address space of a process, or in shared memory environments with multiple processes. Likewise, interprocess communication systems do not currently allow messages to be sent directly between entities within a process. The objective of this project is to build a system which combines standard interfaces for lightweight threads, pthreads, and interprocess communication, MPI, to support point-to-point communication between any two threads in a distributed memory system.

The Chant runtime system has been built using four layers: point-to-point communication, remote service requests, remote thread operations, and collective communication for thread groups. In the last few months we have been exploring extensions which will support load-balancing of distributed computations via migration of light-weight threads. The major issue is the migration of threads in the presence of pointers. We have designed a system which allows threads which use

pointers to be moved from one processor to another within a homogeneous system. The runtime system keeps track of the thread heap along with all the pointers (both stack and heap) used by the thread. The latter requires users to register each of the pointers used by the thread. When migrating a thread, the heap is transferred to the destination processor along with the thread stack. The major overhead is on the destination processor which then has to update all the pointers to point to the new locations based on the offsets calculated from the old locations.

We are in the process of implementing the system. There are some shortcomings of the current system. In particular, we are still exploring how to handle shared data, i.e., when a pointer points to global address space or into another stack's heap area.

This work is being done in collaboration with Matthew Haines of the University of Wyoming and David Cronk, a VILaP graduate student.

Evaluation of HPF

The stated goal of High Performance Fortran (HPF) was to "address the problems of writing data parallel programs where the distribution of data affects performance." We have been using data parallel codes of interest to NASA to evaluate the effectiveness of the language features of HPF.

We have been comparing the performance of three different compilers: APR's xHPF, PGI's pghpf and IBM's xlHPF on the IBM SP-2. We have used a number of CFD related codes ranging from simple one-dimensional problems to a complex multiblock application, TLNS3D. Our experiments have shown that simple single grid-based codes with nearest neighbor communication are easily handled by the current HPF compilers. The performance is obviously dependent on the size of the problem and levels off as the overhead of communication overtakes the amount of computation per processor. For more complex problems, the code has to be expressed "properly" so that the compiler has enough information to generate efficient code.

The multi-block TLNS3D is a more complex code in which different levels of parallelism can be exploited. We have experimented with a model version of the code to exploit the inter-block parallelism. (An earlier version exploited intra-block parallelism). The performance of the communication generated by the compilers is comparable to hand-coded MPI-based version of the code but, due to compiler transformations, the node code has very poor comparative performance.

A new version of HPF (HPF 2.0) has just been released for comments which includes several extensions resulting from our experiments. We plan to continue the evaluation of the new features of HPF as they are incorporated in the compilers.

This work is being done in collaboration with Kevin Roe, a VILaP graduate student.

SmartFiles: A mechanism for data interoperability

Data files for scientific and engineering codes typically consist of a series of raw data values whose description is buried in the programs that interact with these files. In this situation, making even minor changes in the file structure or sharing files between programs (interoperability) can only be done after careful examination of the data files and the I/O statements of the programs interacting with this file. In short, scientific data files lack self-describing capabilities that would improve their interoperability. By applying an object-oriented methodology to data files, we can add the intelligence required to improve data interoperability and provide an elegant mechanism

for supporting complex, evolving, or multidisciplinary applications, while still supporting legacy codes and data files. As a result, scientists and engineers can share datasets with far greater ease, simplifying multidisciplinary applications and facilitating remote collaboration between scientists.

Our approach to the software engineering issues inherent in data files is to “smart files,” to replace the current “dumb” data files. Our goal is to apply the principles of encapsulation, modularity, and inheritance to data files, resulting in a cleaner file abstraction and greatly simplifying the interaction of users with complex scientific and engineering programs.

In the last year we have extended our initial prototype with full functionality as envisioned in our original report, and have evaluated the system in comparison to normal UNIX file I/O routines. We are in the process of completing a journal paper that summarizes the current system.

We plan to advertise the system for use by the scientific computing community in the hopes that it will ease the burden of working with multiple data files stored in different formats. Based on the real-world experiences of scientists, we will modify the system to suit their needs.

This research was conducted in collaboration with Matthew Haines of the University of Wyoming and John Van Rosendale of ICASE.

DAVID NICOL

RITE (Reliability Interface Tool Extension)

One of the principle reasons reliability analysis is treated as an afterthought in system design is that the tools and modeling methodology used is very different from that used to design the functional behavior. Our objective is to address the problem by integrating the reliability analysis directly into the functional design tool.

Our approach is to develop a reliability analysis engine designed specifically for such integration. Our current effort—RITE—is an enhancement of our previous work with the tool REST. RITE addresses several shortcomings of REST, specifically by using new mathematics for the analysis that correct REST’s breakdown on long mission times, and by analyzing multiple mission times concurrently. We have built this engine, tested its acceleration over REST, and enhanced its simulation-based analysis option.

Future work includes creating a distributed version of RITE that exploits the inherent parallelism in the analysis.

Automated parallelization of discrete state-space generation and analysis

A rich variety of modeling and analysis paradigms involve generating a discrete state-space, and then analyzing that space to determine performance and reliability measures. Unfortunately these state spaces can grow quite large, limiting the applicability of a discrete state-space analysis approach. Our objective is to increase the size of models that can be so approached by using multiple processors to accelerate the generation and solution time, and to make available the larger memory of the distributed system.

Our approach is to develop a technique for automatically mapping states to processors, on-the-fly as they are generated, with low overhead and capability of remapping. We accomplished this by: (1) modifying the former state-hashing approach to: (a) generate automatically a way of classifying states based on their position in an AVL tree; and (b) map statically defined state

classes dynamically to processors, as needed. (2) Demonstrating that automated approach actually outperforms former hand-tuned approach on large system models; (3) Demonstrating nearly linear speedups on the SP2; and (4) Implementing block Gauss-Seidel solver.

In the future we plan to investigate performance optimization, and to explore ways of automatically "splitting" a class, to deal with unpartitionable monster classes

This research was conducted in collaboration with Gianfranco Ciardo of The College of William and Mary.

CAN ÖZTURAN

Distributed environment for unstructured meshes

Parallel adaptive PDE solvers on two- and three-dimensional complex geometry need a greater repertoire of mesh entity adjacency relationships for h-refinement, parallel refinement, coarsening and dynamic load balancing routines. The implementation of these routines need specialized data structures which provide fast updates not only on a single processor but also during entity migrations among processors. The objective of this project is to develop an environment, Parallel Mesh Environment (PME), for distributed unstructured grid manipulations.

Hierarchical mesh entity data structures are used which store local adjacencies of entities of consecutive dimensional order, i.e., in the order of region, face, edge and vertex. The mesh is distributed by elements which has the effect of duplicating the lower level entities on partition boundaries. For load balancing, PME implements parallel moment of inertia geometric partitioner. Parallel sample sort is used in the partitioner and the mesh is migrated after load balancing by using the owner updates heuristic. Two-dimensional mesh refinement is carried out by parallelizing Rivara's longest edge bisection algorithm.

PME is implemented using C and the MPI message passing libraries. The software has been tested on a variety of platforms such as networks of SUN and SGI workstations, IBM SP2, Paragon and also on a Pentium 75, 8 Megabyte memory system running Linux with virtual processors. For example, on a network of SGI workstations, partitioning and distribution of a 85K tetrahedral mesh of Onera-M6 wing onto 2, 4 and 8 processors takes 40, 59 and 71 seconds respectively.

We are currently implementing refinement for 3D tetrahedral meshes and interfacing PME with the Petsc software. Future plans will include development of parallel mesh coarsening algorithms and implementation of hp-adaptive solvers for elliptic problems.

DEBORAH F. PILKEY

Numerical considerations in self-diagnostics using damping

Damage detection and diagnostic techniques using vibration responses that depend on analytical models provide more information about a structure's integrity than those that are not model based. The drawback of these approaches is that some form of a workable model is required. Typically, models of practical structures and their corresponding computational effort are very large. One method of detecting damage in a structure is to measure excess energy dissipation, which can be seen in damping matrices. Calculating damping matrices is important because there is a correspondence between a change in the damping matrix and the change in the health of a

structure. The objective of this research is to investigate the numerical problems associated with computing damping matrices using inverse eigenvalue methods.

The proposed technique requires *a priori* knowledge of the structural mass matrix, which can be obtained using a finite element model. Experimental vibration response is obtained and relayed to transfer-function data, which is then applied to inverse eigenvalue methods. Numerical simulations have been performed on 100-degree-of-freedom models to test the effectiveness of the algorithm and the usefulness of parallel computation for this problem. The work performed shows substantial speedup using High Performance Fortran for the iterative algorithm.

Future plans include investigation and implementation of other parallelization methods for the iterative algorithm tested. Additional damping identification routines will be sought and tested for efficiency in large-scale applications.

This research was conducted in collaboration with D.J. Inman of Virginia Tech.

ALEX POTHEN

A parallel sparse indefinite solver

Large, sparse, symmetric indefinite systems of equations occur in computational structural mechanics, electromagnetics, and linear and nonlinear programming. Our goal is to develop parallel algorithms and software for solving sparse, symmetric indefinite linear equations on distributed-memory multiprocessors.

New parallel algorithms and dynamic data structures were developed to deal with the irregular computation caused by sparsity and numerical pivoting. An exhaustive pivoting strategy suitable for parallelism was used. The multifrontal method was used to organize the factorization. MPI was employed for portability. This is the first parallel solver known to us for the sparse indefinite problem. Extensive testing shows that the solver can accurately solve indefinite systems from structural analysis and linear programming.

We will tune our solver for improving the performance, and reimplement parts of the algorithm as needed. An out-of-core solver will be implemented using MPI-IO. A complex version for Helmholtz problems will be created. We are collaborating with experts in the application areas where this solver would be beneficial.

This is joint work with Yogin Campbell of Old Dominion University.

A microeconomic job scheduler for parallel computers

Scheduling jobs on parallel and cluster computers is a problem that system administrators face, but little is known about how different scheduling policies perform under various job distributions. Our goal is to develop a job scheduler for parallel computers that (i) has robust performance across a broad range of job distributions, and (ii) is flexible in permitting trade-offs between user response times and system utilization.

We extend a microeconomic scheduler that we had developed in earlier work by changing the income rate to be constant, proportional to and inversely proportional to the cumulative computation time of the job. We then systematically design a set of experiments to compare the performance of these scheduling policies against First-Come-First-Served (FCFS), FCFS with reservation, and Shortest-Cumulative-Demand-First policies. Our results show that microeconomic scheduling policies exhibit *robust and flexible performance* across a broad range of parameters. Changing the

income-distribution rule among the jobs of a user is an effective way to trade between the conflicting goals of user and system response times on one hand, and system utilization and maximum waiting time on the other hand.

We are building a simulation tool to compare the performance of these scheduling policies over a range of job distributions. This software will be publicly available, and can be used to evaluate the performance of various scheduling policies on particular job distributions.

This work was done in collaboration with Ion Stoica of Carnegie Mellon University.

Two wavefront reduction ordering algorithms

The wavefront (variable row bandwidth) minimization problem for sparse matrices and graphs has applications in organizing irregular data structures to reduce cache misses, in incomplete factorization preconditioners, and in frontal methods for solving sparse systems of equations. We are developing ordering algorithms and software to reduce the wavefront of sparse matrices and graphs.

A spectral algorithm that relies on an eigenvector of the associated Laplacian matrix was developed. A combinatorial algorithm, the Sloan algorithm, that relies on a two-component priority function, was improved in terms of quality and speed. A hybrid algorithm that combines the two was also developed. The hybrid is currently the best in terms of quality, and Sloan is the fastest among these. We provided an $n \log n + e$ implementation of Sloan; previous best was $n^{3/2}$ for 2D, and $n^{5/3}$ for 3D problems with good separators (here n is number of vertices, e is number of edges). This improved algorithm was practically much faster than earlier ones. We identified two classes of problem behaviors for the Sloan algorithm, and thereby improved our understanding of this algorithm. Three versions of the software (for PETSc, for Matlab, and a stand-alone version) were created and made available to users.

We will look at cache performance of irregular data structures searched in a wavefront reducing ordering. We will consider orderings that consider weights, and apply to preconditioners.

This is joint work with Gary Kumfert of Old Dominion University.

ARUN K. SOMANI

Reliability modeling of structured systems

A large number of systems are implemented using regular interconnected topologies. Markov analysis of such systems results in large state spaces. Our objective in this research is to explore symmetry, in particular rotational and permutational, of such systems to achieve a significant reduction in the size of the state space required to analyze them. The resulting much smaller state spaces allow analyses of very large systems.

To explore symmetry, we define equivalent classes of states and develop an algorithm to generate small state spaces and the corresponding Markov chain for systems with permutation symmetries. In such systems, the relative position of a node is not importance. We exploit permutation and rotational symmetries present in the system to our advantage in state space generation. The state space generation process is also simplified. A complete Markov chain is generated as part of the algorithm. We demonstrate our technique using several examples. Our technique is very useful in the exact analysis of large systems.

We are going to implement this algorithm as part of a analysis tool, called HIMAP (Hierarchical Modeling and Analysis Package) being developed at the University of Washington. This research in part is also being supported by the Boeing Company.

Study of cache memory error propagation

Cache memory is a small, fast, memory system that holds frequently used data. With increasing processor speed, designs follow aggressive practices in the design of cache memories. Such design practices increase the probability of fault occurrence and the presence of latent errors as a processor allows a short duration for read and write. The fault may corrupt the cache memory system or lead to an erroneous internal CPU state. The presence of erroneous data in a processor register or in a cache memory location may corrupt the results due to incorrect data usage, wrong instruction execution, or even the program taking an incorrect branch. This may further corrupt the cache memory locations or lead to an erroneous internal CPU state. The objective of this research is to investigate the error propagation due to transient faults either in the cache memory itself or in the processor's registers or both. The information gained from such an investigation should lead to the development of more effective error recovery mechanisms against failures due to transient faults arising in the machine's cache memory and register set.

We develop and solve a model for error propagation for transient faults in a cache location or a processor register. This model includes the execution of a program and the effect that has on the state of processor instructions, state, and the contents of cache memory. We establish that even though the computer system is capable of recovering about 50% of the time from the effect of a single erroneous cache location/processor register, the other 50% of the time, error recovery is affected only through specific recovery mechanisms. We also verify our results using software injected errors on a real machine running application programs. Our results are obtained using both a discrete-time Markov model and by means of error injection on a real system.

We are currently developing techniques for fault and error detection in cache memories.

This research was conducted in collaboration with Kishor S. Trivedi of Duke University.

Minimizing overhead in parallel algorithms through overlapping communication/computation

One of the major goals in the design of parallel processing machines and algorithms is to reduce the effects of the overhead introduced when a given problem is parallelized. A key contributor to overhead is communication time. Many architectures try to reduce this overhead by minimizing the actual time for a communication to occur, including latency and bandwidth figures. Another approach is to hide communication by overlapping it with computation. The objective of this research is explore techniques to hide communication time at the coarse level by overlapping computing and communication present to achieve higher speed ups. These techniques are easily extended for use in compiler support of parallel programming.

We develop methodology to partition and schedule computations and communication tasks in parallel to exploit overlapping on parallel machines. We assume that the underlying hardware is capable of supporting these mechanisms efficiently. We use Proteus, a parallel reconfigurable machine designed and built at the University of Washington, for our experimental work. We demonstrate that by careful planning and scheduling, the machine is able to exploit full parallelism while hiding the latency of communication and utilize the full bandwidth of communication network.

The net result is that communication system need not be over-designed. We show that the reducing overhead is not a linear function with respect to bandwidth increase. A rather large increase in an effective bandwidth is necessary to decrease the overhead to an acceptable level. Overlapping communications/computation allows for a reduction in the overhead without unacceptably large increases in the effective bandwidth. We also show the complexity, or rather simplicity, in achieving complete exchange on the Proteus Machine using this approach in most efficient manner.

We are currently adopting this methodology to other applications on the Proteus machine.

XIAN-HE SUN

Software environment for non-dedicated network computing

In a non-dedicated network environment, computers are privately owned. Individual owners do not want to see their systems being saturated by others when they need them. This means that privately owned machines may only be used for parallel processing on an "availability" basis. In addition, a non-dedicated network of computers is more likely to be heterogeneous than is a dedicated system. "Availability" and "heterogeneity" are new issues of distributed network computing, which do not arise in tightly coupled parallel systems. Competition for computing resources does not lead to guaranteed high performance. To simultaneously utilize idle machines and maintain high capabilities for local computations, network process migration mechanism is proposed as a solution. The simple idea underlying this mechanism is that when the workload of a distributed system becomes unbalanced, parallel processes residing on overloaded machines are migrated to other available machines.

We have introduced a high-level mechanism and its associated methodologies to support efficient process migration in a non-dedicated, heterogeneous network computing environment. The newly proposed network process-migration mechanism is general. It can be applied to any distributed network environment. In particular, based on this mechanism, a software system named MpPVM is designed and implemented to support efficient process migration for PVM application programs. We have studied the interaction between process migration and resource management and proposed modifications of pvmd and pvmlib of PVM to maintain reliable data communications among processes in a migration environment. Our implementation and experimental results confirm the applicability and potential of the proposed mechanism in a non-dedicated heterogeneous environment. Experimental results indicate the MpPVM software system is efficient and is scalable in the sense that it can carry out a process migration in tenth of milliseconds and the migration costs becomes less notable when the problem and ensemble increase.

The implementation of MpPVM is still preliminary. The Mpd and Mlibpvm need to be improved to maintain high reliability and efficiency on various parallel applications. The MCL software and its supporting tools are also under development. To support an effective process migration environment, an efficient scheduler which controls workload of the distributed environment is inevitable. We plan to develop a scheduler that can efficiently exploit process migrations in a non-dedicated, heterogeneous computing environment. The performance evaluation and prediction according to effects of process migrations on various algorithms and heterogeneous platforms will also be investigated.

KISHOR TRIVEDI

Modeling and simulation of fluid stochastic Petri nets

One of the difficulties of existing modeling tools is that the state-spaces of the systems they analyze grow too large. One way of approaching this problem is to use continuous quantities to approximate discrete ones, e.g., use fluid flow to model communication traffic. The objective of this research is to explore the possibilities and potentials of this approach by augmenting the widely used Petri net paradigm with continuous "fluid" constructs, making Fluid Stochastic Petri Nets (FSPN).

FSPNs is a novel paradigm that can permit stochastic modeling of systems with mixed continuous and discrete state spaces. The formalism allows a concise specification and automated solution of such systems. FSPNs can be used to approximate discrete state space systems with very large state spaces as well. In fact, in this context, they can be directly applied to the performance analysis of high speed communication networks.

We have brought together computer scientists with applied mathematicians to approach the numerical problems associated with solving FSPN models. We have demonstrated the utility of the approach, showing that on a model of an ATM switch a FSPN formulation solves the model nearly 50 times faster than its equivalent discrete counterpart, with quantitatively indistinguishable results. We have begun to explore the issues involved in mixed continuous-discrete simulation of FSPN for the large models for which closed analysis is intractable.

We are designing and will build a FSPN simulation engine, and integrate it with a graphical system design tool under development at Duke University. We are in discussion with several NASA researchers in finding interesting applications in the domain of real-time control where FSPN may provide a new way of studying the behavior of these systems.

This research was conducted in collaboration with David Nicol of Dartmouth College and Gianfranco Ciardo of The College of William and Mary, and Graham Horton of Universitat Erlangen-Nurmburg, Germany.

LINDA F. WILSON

Automated load balancing in parallel discrete-event simulation

Discrete-event simulation is a powerful, versatile tool that can be used to examine a variety of performance-related issues such as the performance of a parallel algorithm or the average aircraft delay in the National Airspace System. Parallel discrete-event simulation offers the potential for significant speedup over sequential simulation. However, high performance is often achieved only after rigorous fine-tuning is used to obtain an efficient mapping of tasks to processors. This research seeks to develop tools for automating the tuning process with little or no modification to the user's simulation code. In particular, we are examining methods to efficiently accomplish automated load balancing in distributed object-oriented simulation tools.

We are using the SPEEDES simulation package to examine performance problems in a variety of large simulation projects. In particular, we are examining the related issues of automated load balancing, object migration, and communication/synchronization in an effort to improve the performance of such tools. Using the Intel Paragon, we modified SPEEDES to perform automated

load balancing using three mapping algorithms that used run-time measurements. We have begun modifications to SPEEDES to perform object migration, which is needed to reach the goal of automated dynamic load balancing. As part of this work, we are porting SPEEDES to use the MPI message-passing library on the IBM SP2.

Once SPEEDES has the ability to migrate objects, run-time measurements of a simulation workload will be used to guide an automated remapping mechanism.

This research was conducted in collaboration with David M. Nicol of Dartmouth College.

REPORTS AND ABSTRACTS

Choudhari, Meelan, and Peter W. Duck: *Nonlinear excitation of inviscid stationary vortex in a boundary-layer flow*. ICASE Report No. 96-27, (NASA CR-198327), April 19, 1996, 18 pages. To be submitted to the Journal of Fluid Mechanics.

Arian, Eyal, and Shlomo Ta'asan: *Analysis of the Hessian for aerodynamic optimization: Inviscid flow*. ICASE Report No. 96-28, (NASA CR-198328), April 29, 1996, 22 pages. To be submitted to the Journal of Computational Physics.

Horton, Graham, and Kishor S. Trivedi: *Computation of the distribution of accumulated reward with fluid stochastic Petri-nets*. ICASE Report No. 96-29, (NASA CR-198330), April 23, 1996, 12 pages. Submitted to the International Performance and Dependability Symposium.

Jameson, Leland: *A wavelet-optimized, very high order adaptive grid and order numerical method*. ICASE Report No. 96-30, (NASA CR-198331), May 3, 1996, 42 pages. To be submitted to SIAM Scientific Computations.

Erlebacher, Gordon, M.Y. Hussaini, and Chi-Wang Shu: *Interaction of a shock with a longitudinal vortex*. ICASE Report No. 96-31, (NASA CR-198332), May 3, 1996, 43 pages. Submitted to the Journal of Fluid Mechanics.

Stoica, Ion, and Alex Pothen: *A robust and flexible microeconomic scheduler for parallel computers*. ICASE Report No. 96-32, (NASA CR-198333), May 3, 1996, 12 pages. Submitted to the Proceedings of the Third International High Performance Computing Conference.

Kozusko, F., D.G. Lasseigne, C.E. Grosch, and T.L. Jackson: *The stability of compressible mixing layers in binary gases*. ICASE Report No. 96-33, (NASA CR-198334), May 14, 1996, 28 pages. To appear in Physics of Fluids.

Stoica, Ion, Florin Sultan, and David Keyes: *Evaluating the hyperbolic model on a variety of architectures*. ICASE Report No. 96-34, (NASA CR-198335), May 14, 1996, 20 pages. To appear in the Proceedings of Europar '96.

Ciardo, Gianfranco, and Marco Tilgner: *The use of Kronecker operators for the solution of generalized stochastic Petri nets*. ICASE Report No. 96-35, (NASA CR-198336), May 28, 1996, 32 pages. Submitted to the Journal of the Association for Computing Machinery.

Kozusko, F., C.E. Grosch, T.J. Jackson, Christopher A. Kennedy, and Thomas B. Gatski: *The structure of variable property, compressible mixing layers in binary gas mixtures*. ICASE Report No. 96-36, (NASA CR-198337), May 28, 1996, 25 pages. To appear in Physics of Fluids.

Le Tallec, Patrick, and Moulay D. Tidriri: *Convergence analysis of domain decomposition algorithms with full overlapping for the advection-diffusion problems*. ICASE Report No. 96-37, (NASA CR-198339), July 25, 1996, 31 pages. Submitted to Mathematics of Computations.

Girimaji, Sharath S.: *A Galilean invariant explicit algebraic Reynolds stress model for curved flows*. ICASE Report No. 96-38, (NASA CR-198340), June 13, 1996, 28 pages. To be submitted to Physics of Fluids.

Cai, Xiao-Chuan, William D. Gropp, David E. Keyes, Robin G. Melvin, and David P. Young: *Parallel Newton-Krylov-Schwarz algorithms for the transonic full potential equation*. ICASE Report No. 96-39, (NASA CR-198341), June 10, 1996, 27 pages. Proceedings of the 1996 Copper Mountain Conference on Iterative Methods; submitted to SIAM Journal of Scientific Computing.

Sorensen, Danny C.: *Implicitly restarted Arnoldi/Lanczos methods for large scale eigenvalue calculations*. ICASE Report No. 96-40, (NASA CR-198342), June 10, 1996, 36 pages. To appear in the Proceedings of the ICASE Workshop on Parallel Numerical Algorithms.

Martin, J.E., and E. Meiburg: *Nonlinear axisymmetric and three-dimensional vorticity dynamics in a swirling jet model*. ICASE Report No. 96-41, (NASA CR-198343), June 10, 1996, 30 pages. Submitted to Physics of Fluids A.

Martinez, D.O., S. Chen, G.D. Doolen, R.H. Kraichnan, L.-P. Wang, and Y. Zhou: *Energy spectrum in the dissipation range of fluid turbulence*. ICASE Report No. 96-42, (NASA CR-198344), June 10, 1996, 14 pages. Submitted to the Journal of Plasma Physics.

Borggaard, Jeff, and John Burns: *A PDE sensitive equation method for optimal aerodynamic design*. ICASE Report No. 96-44, (NASA CR-198349), June 24, 1996, 40 pages. Submitted to the Journal of Computational Physics.

Le Tallec, Patrick, and Moulay D. Tidriri: *Maximum principles and application to the analysis of an explicit time marching algorithm*. ICASE Report No. 96-45, (NASA CR-201584), August 13, 1996, 33 pages. To be submitted to SIAM Journal on Mathematical Analysis.

Kelley, C.T., and David E. Keyes: *Convergence analysis of pseudo-transient continuation*. ICASE Report No. 96-46, (NASA CR-201585), July 19, 1996, 21 pages. Submitted to the SIAM Journal of Numerical Analysis.

Abarbanel, Saul, and Adi Ditkowski: *Multi-dimensional asymptotically stable finite difference schemes for the advection-diffusion equation*. ICASE Report No. 96-47, (NASA CR-201586), July 26, 1996, 36 pages. To be submitted to the Journal of Computational Physics.

Cai, Xiao-Chuan, Charbel Farhat, and Marcus Sarkis: *Variable degree Schwarz methods for the implicit solution of unsteady compressible Navier-Stokes equations on two-dimensional unstructured meshes*. ICASE Report No. 96-48, (NASA CR-201588), August 13, 1996, 21 pages. Submitted to SIAM Journal of Scientific Computing.

Ristorcelli, J.R., and G.A. Blaisdell: *Consistent initial conditions for the DNS of compressible turbulence*. ICASE Report No. 96-49, (NASA CR-201589), July 26, 1996, 9 pages. To be submitted to Physics of Fluids A.

Gottlieb, Sigal, and Chi-Wang Shu: *Total variation diminishing Runge-Kutta schemes*. ICASE Report No. 96-50, (NASA CR-201591), July 31, 1996, 20 pages. Submitted to Mathematics of Computation.

Atkins, H.L., and Chi-Wang Shu: *Quadrature-free implementation of discontinuous Galerkin method for hyperbolic equations*. ICASE Report No. 96-51, (NASA CR-201594), August 20, 1996, 35 pages. Submitted to AIAA Journal.

Rubinstein, Robert, and Gordon Erlebacher: *Transport coefficients in weakly compressible turbulence*. ICASE Report No. 96-52, (NASA CR-201595), August 20, 1996, 47 pages. Submitted to Physics of Fluids.

Ta'asan, Shlomo, and Hong Zhang: *Fourier-Laplace analysis of multigrid waveform relaxation method for hyperbolic equations*. ICASE Report No. 96-53, (NASA CR-201598), August 30, 1996, 15 pages. To appear in BIT Numerical Mathematics.

Bokhari, Shahid, and David M. Nicol: *Balancing contention and synchronization on the Intel Paragon*. ICASE Report No. 96-54, (NASA CR-201599), August 30, 1996, 22 pages. Submitted to IEEE Parallel and Distributed Technology.

Jordan, Harry F., Kyungsook Y. Lee, Rajgopal Kannan, Coke Reed, Jon R. Sauer, and Douglas E. Straub: *A high speed optical radix sorter*. ICASE Report No. 96-55, (NASA CR-201600), August 30, 1996, 18 pages. Submitted to the Third International Symposium on High-Performance Computer Architecture.

ICASE COLLOQUIA
April 1, 1996 – September 30, 1996

Name/Affiliation/Title	Date
Taylor, Arthur, Old Dominion University "Observations on Discrete Second-Order Aerodynamic Sensitivity Analysis"	April 5
Ben-Tal, Aharon, TECHNION, Israel "Layout Optimization of Engineering Structures by Mathematical Programming Methods"	April 9
Saied, Faisal, University of Illinois at Urbana-Champaign "Multiscale Solvers for Computational Science"	April 24
Bruun, Hans, Bradford University, UK "Hot-Wire Anemometry"	April 25
Bruun, Hans, Bradford University, UK "Flying Hot-Wire Probe Study of Flow over Model Aerofoils and Comparison with CFD Evaluations"	April 26
Cormen, Thomas, Dartmouth College "ViC*: A Compiler for Virtual-Memory C**"	April 29
Lilley, Geoffrey, University of Southampton, UK "Aeroacoustic Theories and the Prediction of Jet Noise"	May 8
Erlebacher, Gordon, ICASE "On the Nonlinear Interactions of Shocks with Organized Structures"	May 17
Nowaczyk, Ron, Clemson University "The Human Factor in Multidisciplinary Design and Analysis"	May 20
Laney, Culbert, University of Colorado at Boulder "Acoustic Effects on Aerodynamic Lift and Drag"	May 24
Burdakov, Oleg, CERFACS, Toulouse, France "Interpolation Methods for Optimization and Nonlinear Equations"	May 24
Banks, David, Mississippi State University "Evenly-Spaced Streamlines"	May 29

Name/Affiliation/Title	Date
Pevchin, Sergei, Virginia Polytechnic Institute and State University "Capture of Contact Surfaces and Shock Waves Using a Discontinuity Confinement Procedure"	May 31
Xu, Kun, Princeton University "Gas-Kinetic Scheme for the Compressible Flow Simulations"	June 6
Banks, David, Mississippi State University "Triangulating Contours of Implicit Functions"	June 7
Mahalov, Alex, Arizona State University "Phase Turbulence in 3D Rotating and Stratified Geophysical Flows"	June 13
Jameson, Antony, Princeton University "Evolution and Current Developments in Automatic Design by Control Theory"	June 14
Sritharan, S., University of Colorado-Boulder and US Navy-San Diego "New Advances in Optimal Control of Aerothermodynamics"	June 17
Mikata, Yozo, Old Dominion University "Micromechanics: Coated Fiber Composites and Flaw Detection"	June 18
Mistree, Farrokh, Georgia Institute of Technology "Integrated Product and Process Design: A Decision-Based Approach"	June 19
Mistree, Farrokh, Georgia Institute of Technology "Living-Systems Analogies in Decision-Based Design"	June 20
Sritharan, S., University of Colorado-Boulder and US Navy-San Diego "Optimal Control of Viscous Flow I – Deterministic Control"	June 25
Domaradzki, Julian, University of Southern California "Non-Eddy Viscosity Subgrid-Scale Model for Turbulent Flows"	June 27
O'Sullivan, Peter, Southern Methodist University, Dallas, Texas "Numerical Simulations of Transitional Pipe Flow and a Viscous Incompressible Round Jet"	July 1
Sritharan, S., University of Colorado-Boulder and US Navy-San Diego "Optimal Control of Viscous Flow II: (Robust) H-Infinity Control"	July 3
Sritharan, S., University of Colorado-Boulder and US Navy-San Diego "Optimal Control of Viscous Flow III: Stochastic Control and Nonlinear Kalman Filters"	July 9

Name/Affiliation/Title	Date
Ryzhov, Oleg, Rensselaer Polytechnic Institute "The Development of Nonlinear Solitary-Like Waves in Boundary Layers and Their Randomization"	July 15
Sanders, William, University of Illinois, Center for Reliable and High Performance Computing "Uniformization Methods in Markov Modeling"	July 17
Das, Indraneel, Rice University "Normal Boundary Intersection: An Alternate Approach to Multicriteria Optimization"	July 19
Mehra, Pankaj, Indian Institute of Technology "Experimental and Analytical Studies in Parallel and Distributed Systems"	July 26
Shu, Chi-Wang, Brown University "Recent Developments on Discontinuous Galerkin Method"	July 26
Berkooz, Gal, BEAM Technologies, Inc. "PDESolve: A Tool for Rapid Development of PDE Simulations"	July 30
Geer, James, State University of New York "Accurate Approximations to a Function and its Derivatives Using Fourier Series Partial"	August 20
Braun, Robert, Vehicles Analysis Branch, NASA LaRC "Collaborative Optimization: An Architecture for Large-Scale Distributed Design"	August 23
Silver, Deborah, Rutgers University "Feature Extractions and Tracking"	August 26
Mitra, Raj, Pennsylvania State University "Overview in CEM and Antenna Design at the Electromagnetics Communication Laboratory"	August 26
Mitter, Sanjoy K., Massachusetts Institute of Technology "Control at Interconnection"	August 27
Delfour, Michel, Universite de Montreal "Shape Optimization and Control"	August 28
Mitter, Sanjoy K., Massachusetts Institute of Technology "Global Stochastic Recursive Algorithms"	August 29

Name/Affiliation/Title	Date
Luo, Li-Shi, Los Alamos National Laboratory "Lattice-Gas Automata and Lattice Boltzmann Methods: New Developments"	September 4
Ozturan, Can, ICASE "Distributed Environment for Unstructured Grids"	September 13
Bernard, Peter, University of Maryland "Transport Physics and its Application to Turbulent Flow Modeling"	September 19
Crockett, Thomas, ICASE "Beyond the Renderer: Software Architecture for Parallel Graphics and Visualization"	September 20

ICASE SUMMER ACTIVITIES

The summer program for 1996 included the following visitors:

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Abarbanel, Saul <i>Applied & Numerical Math</i>	Brown University	7/08 – 8/01 8/18 – 8/23
Agarwal, Ramesh <i>Applied & Numerical Math</i>	Wichita State University	8/05 – 8/09
Baggag, Abdelkader <i>Computer Science</i>	University of Minnesota	6/17 – 9/13
Balachandar, Sivaramakrishnan <i>Applied & Numerical Math</i>	University of Illinois	7/01 – 7/12
Banks, David <i>Computer Science</i>	Mississippi State University	5/20 – 6/14
Banks, H. Thomas <i>Applied & Numerical Math</i>	North Carolina State University	5/06 – 5/10 8/21 – 8/23 8/28 – 8/30 9/12 – 9/13 9/25 – 9/27
Bataille, Francoise <i>Fluid Mechanics</i>	Institut National des Sciences Appliquées de Lyon, France	7/15 – 8/14
Bayliss, Alvin <i>Fluid Mechanics</i>	Northwestern University	5/01 – 5/03 5/20 – 5/23 6/17 – 6/20 6/30 – 7/03 8/05 – 8/08 8/20 – 8/23 9/09 – 9/12
Berger, Stanley A. <i>Fluid Mechanics</i>	University of California, Berkeley	8/05 – 8/16
Bloebaum, Christina L. <i>Applied & Numerical Math</i>	State University of New York at Buffalo	8/19 – 8/23

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Boersma, Bendicks Jan <i>Fluid Mechanics</i>	Delft University of Technology	7/15 – 8/09
Bokhari, Shahid <i>Computer Science</i>	University of Engineering and Technology, Lahore, Pakistan	6/07 – 9/07
Booker, Andrew <i>Applied & Numerical Math</i>	The Boeing Company	8/12 – 8/16
Booty, Michael <i>Applied & Numerical Math</i>	New Jersey Institute of Technology	5/13 – 5/24
Brandt, Achi <i>Applied & Numerical Math</i>	Weizmann Institute of Technology	7/01 – 7/03
Burns, John <i>Applied & Numerical Math</i>	Virginia Polytechnic Institute and State University	8/12 – 8/23
Cai, Xiao-Chuan <i>Computer Science</i>	University of Colorado	5/16 – 5/22 8/12 – 8/23
Caromel, Denis R. <i>Computer Science</i>	Universite de Nice-Sophia Antipolis, France	7/29 – 8/02
Chan, Tony <i>Applied & Numerical Math</i>	University of California, Los Angeles	7/29 – 8/02
Chapman, Barbara <i>Computer Science</i>	University of Vienna, Austria	7/29 – 8/09
Chrisochoides, Nikos <i>Computer Science</i>	Cornell University	7/29 – 8/02
Ciardo, Gianfranco <i>Computer Science</i>	The College of William & Mary	7/15 – 8/16
Criminale, William O. <i>Fluid Mechanics</i>	University of Washington, Seattle	6/17 – 7/05
Das, Indraneel <i>Applied & Numerical Math</i>	Rice University	5/20 – 7/26
Ditkowski, Adi <i>Applied & Numerical Math</i>	Tel-Aviv University, Israel	7/08 – 8/02

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Domaradzki, Julian A. <i>Fluid Mechanics</i>	University of Southern California	6/03 – 6/28
Follett, William <i>Applied & Numerical Math</i>	Rocketdyne International	8/05 – 8/16
Fox, Geoffrey <i>Computer Science</i>	Syracuse University	6/10 – 6/14
Funaro, Daniele <i>Applied & Numerical Math</i>	Universita di Modena, Italy	7/29 – 8/16
Gannon, Dennis <i>Computer Science</i>	Indiana University at Bloomington	7/31 – 8/02
Geer, James <i>Applied & Numerical Math</i>	State University of New York	8/19 – 8/23
Gottlieb, David <i>Applied & Numerical Math</i>	Brown University	7/08 – 7/19
Grosch, Chester E. <i>Fluid Mechanics</i>	Old Dominion University	5/06 – 5/31
Gunzburger, Max <i>Applied & Numerical Math</i>	Iowa State University	5/28 – 5/31
Hafez, Mohamed <i>Applied & Numerical Math</i>	University of California, Davis	8/05 – 8/09
Haines, Matthew <i>Computer Science</i>	University of Wyoming	5/13 – 6/14 7/29 – 8/02
Hittinger, Jeffrey A. <i>Computer Science</i>	The University of Michigan, Ann Arbor	5/23 – 8/16
Holt, Maurice <i>Applied & Numerical Math</i>	University of California, Berkeley	7/03 – 7/10
Horton, Graham <i>Computer Science</i>	Universitat Erlangen-Nurmburg, Germany	8/11 – 8/28
Hu, Fang <i>Fluid Mechanics</i>	Old Dominion University	5/13 – 7/19

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Jones, Don <i>Applied & Numerical Math</i>	GM Research & Development	8/12 – 8/16
Jordan, Harry <i>Computer Science</i>	University of Colorado, Boulder	6/17 – 6/28
Kapila, Ashwani <i>Fluid Mechanics</i>	Rensselaer Polytechnic Institute	8/12 – 8/16
Kesselman, Carl <i>Computer Science</i>	California Institute of Technology	7/29 – 8/02
Keyes, David E. <i>Computer Science</i>	Old Dominion University	5/20 – 5/31 7/01 – 8/30
Koehler, James <i>Applied & Numerical Math</i>	University of Colorado, Denver	8/12 – 8/15
Kopriva, David <i>Applied & Numerical Math</i>	Florida State University	5/06 – 5/10
Kumar, Vipin <i>Computer Science</i>	University of Minnesota	7/29 – 8/09
Langendoen, Koen G. <i>Computer Science</i>	Vrije Universiteit, Holland	5/28 – 5/31
Leutenegger, Scott <i>Computer Science</i>	University of Denver	7/24 – 8/18
Lohner, Rainald <i>Computer Science</i>	George Mason University	8/05 – 8/09
MacCormack, Robert <i>Applied & Numerical Math</i>	Stanford University	8/05 – 8/09
Mahalov, Alex <i>Fluid Mechanics</i>	Arizona State University	5/28 – 6/14
McCormick, Stephen <i>Applied & Numerical Math</i>	University of Colorado, Boulder	5/06 – 5/09
Mitter, Sanjoy <i>Applied & Numerical Math</i>	Massachusetts Institute of Technology	8/26 – 8/30

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Mukherjee, Amar <i>Computer Science</i>	University of Central Florida	7/01 – 7/12
Nicol, David <i>Computer Science</i>	The College of William & Mary	5/01 – 5/15 9/23 – 9/27
Nicolaides, R.A. <i>Applied & Numerical Math</i>	Carnegie Mellon University	5/06 – 5/08
Papageorgiou, Demetrius <i>Fluid Mechanics</i>	New Jersey Institute of Technology	5/13 – 5/24
Patera, Anthony <i>Applied & Numerical Math</i>	Massachusetts Institute of Technology	8/12 – 8/14
Pavarino, Luca <i>Applied & Numerical Math</i>	Universita di Pavia, Italy	8/19 – 9/13
Pilkey, Deborah <i>Computer Science</i>	Virginia Polytechnic Institute and State University	5/20 – 8/16
Pothen, Alex <i>Computer Science</i>	Old Dominion University	5/06 – 8/31
Reif, John <i>Computer Science</i>	Duke University	6/05 – 6/07 7/15 – 7/19 8/19 – 8/21
Roe, Philip <i>Fluid Mechanics</i>	University of Michigan	7/01 – 8/09
Saad, Yousef <i>Computer Science</i>	University of Minnesota	8/05 – 8/09
Shu, Chi-Wang <i>Applied & Numerical Math</i>	Brown University	5/13 – 5/31 7/22 – 7/26
Smith, Ralph <i>Applied & Numerical Math</i>	Iowa State University	6/03 – 6/07 8/19 – 8/23
Somani, Arun K. <i>Computer Science</i>	University of Washington, Seattle	8/19 – 8/30
Speziale, Charles <i>Fluid Mechanics</i>	Boston University	8/12 – 8/16

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Sritharan, S.S. <i>Fluid Mechanics</i>	University of Colorado, Boulder	6/17 – 7/19
Sun, Xian-He <i>Computer Science</i>	Louisiana State University	7/22 – 8/09
Ta'asan, Shlomo <i>Applied & Numerical Math</i>	Carnegie Mellon University	5/27 – 6/08 8/05 – 8/09
Thangam, Siva <i>Fluid Mechanics</i>	Stevens Institute of Technology	6/10 – 7/05
Ting, Lu <i>Fluid Mechanics</i>	New York University	5/20 – 5/24 6/10 – 6/14 8/12 – 8/23
Torczon, Virginia <i>Computer Science</i>	The College of William & Mary	5/17, 5/29, 5/31 8/19 – 8/23
Trivedi, Kishor <i>Computer Science</i>	Duke University	8/12 – 8/23
Trosset, Michael <i>Applied & Numerical Math</i>	University of Arizona	8/09 – 8/16
Turkel, Eli <i>Applied & Numerical Math</i>	Tel-Aviv University, Israel	7/01 – 8/09
van Leer, Bram <i>Applied & Numerical Math</i>	The University of Michigan, Ann Arbor	7/15 – 8/09
Watson, Layne <i>Applied & Numerical Math</i>	Virginia Polytechnic Institute and State University	8/12 – 8/16
Widlund, Olof B. <i>Computer Science</i>	New York University	8/19 – 8/23
Woodruff, Stephen L. <i>Fluid Mechanics</i>	Brown University	5/28 – 6/07
Xu, Kun <i>Applied & Numerical Math</i>	Princeton University	6/05 – 6/08
Yang, Chuan-Kai <i>Computer Science</i>	SUNY at Stony Brook	6/17 – 8/13

<u>VISITOR and AREA OF RESEARCH</u>	<u>AFFILIATION</u>	<u>DATE OF VISIT</u>
Yeung, Pui-Kuen <i>Fluid Mechanics</i>	Georgia Institute of Technology	6/17 – 6/28
Younis, Bassom A. <i>Fluid Mechanics</i>	City University of London, England	6/01 – 8/31
Zhang-Sun, Hong <i>Applied & Numerical Math</i>	Clemson University	7/29 – 8/09
Zhuang, Mei <i>Fluid Mechanics</i>	Michigan State University	8/05 – 8/16
Zima, Hans P. <i>Computer Science</i>	University of Vienna, Austria	7/29 – 8/23

OTHER ACTIVITIES

On May 29–31, 1996, ICASE and NASA LaRC co-sponsored a **Workshop on Computational Electromagnetics** at the OMNI Hotel in Newport News, VA. The objective of the Workshop was to show a broad cross-section of the tools and applications of computational electromagnetics (CEM). There were 63 attendees and a formal proceedings will be published.

On June 10–13, 1996, ICASE and NASA LaRC co-sponsored a short course on “**Introduction to Emerging Web Technologies for Scientific Computations**” taught by Professor Geoffrey Fox, Syracuse University. This was followed by a one day workshop on June 14 in which several invited speakers discussed their own activities in the scientific use of Web Technologies. Eighty NASA scientists and contractors attended this week-long course.

On July 29–August 2, 1996, ICASE and NASA LaRC co-sponsored a **Workshop on Programming Models for Future Parallel Architectures**. The objective of this workshop was to discuss the structure and form of a uniform programming model for future generations of parallel architectures. There were 20 researchers from universities, national labs, and industry in attendance.

On August 5–7, 1996, ICASE and NASA LaRC co-sponsored a **Workshop on Barriers and Challenges in Computational Fluid Dynamics** at the Holiday Inn in Hampton, VA. The objective of the Workshop was to catalog the barriers and challenges in CFD and to identify new application areas that could benefit from CFD techniques. There were 67 attendees, and a formal proceedings will be published.

On August 12–16, 1996, ICASE and NASA LaRC co-sponsored a **Workshop on Approximation Methods**. The objective of this workshop was to identify significant research areas in the systematic algorithmic use of approximations and surrogates to reduce computational cost in engineering design and optimization. There were a dozen participants from academia, industry, and NASA. A written report of the meeting is in preparation.

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970. Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Etta M. Blair, Accounting Supervisor

Barbara A. Cardasis, Administrative Secretary

Shelly M. Johnson, Executive Secretary/Visitor Coordinator

Shannon L. Keeter, Technical Publications Secretary

Rachel A. Lomas, Payroll and Accounting Clerk

Emily N. Todd, Conference Manager

Gwendolyn W. Wesson, Contract Accounting Clerk

Leon M. Clancy, Senior System Manager

Bryan K. Hess, Assistant System Manager

Gregory P. Wheeler, System Operator

II. SCIENCE COUNCIL

Ivo Babuska, Robert Trull Chair in Engineering, The University of Texas-Austin.

Geoffrey Fox, Director, Northeast Parallel Architectural Center, Syracuse University.

Dennis Gannon, Professor, Center for Innovative Computer Applications, Indiana University.

Ashwani Kapila, Professor, Department of Mathematics and Science, Rensselaer Polytechnic Institute.

James P. Kendall, Jet Propulsion Laboratory.

Heinz-Otto Kreiss, Professor, Department of Mathematics, University of California at Los Angeles.

Sanjoy Mitter, Professor of Electrical Engineering, Massachusetts Institute of Technology.

Steven A. Orszag, Professor, Program in Applied & Computational Mathematics, Princeton University.

Paul Rubbert, Unit Chief, Boeing Commercial Airplane Group.

Ahmed Sameh, Department Head of Computer Science, University of Minnesota.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Gordon Erlebacher - Ph.D., Plasma Physics, Columbia University, 1983. Fluid Mechanics [Transition and Turbulence]. (November 1989 to August 1996)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 1999)

IV. SENIOR STAFF SCIENTISTS

Sharath S. Girimaji - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1990. Fluid Mechanics [Turbulence and Combustion]. (July 1993 to August 1997)

Thomas Jackson - Ph.D., Mathematics, Rensselaer Polytechnic Institute, 1985. Fluid Mechanics. (January 1994 to January 1997)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to April 1998)

Josip Loncaric - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to February 1997)

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1987 to August 1997)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Turbulence Modeling]. (January 1995 to January 1998)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to October 1996)

V. Venkatakrisnan - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1987. Applied & Numerical Mathematics [Computational Aerodynamics]. (June 1993 to September 1997)

Ye Zhou - Ph.D., Physics, College of William and Mary, 1987. Fluid Mechanics [Turbulence Modeling]. (October 1992 to September 1996)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to January 1998)

Eyal Arian - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1995. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (October 1994 to September 1996)

Phillip M. Dickens - Ph.D., Computer Science, University of Virginia, 1992. Computer Science [System Software]. (January 1993 to August 1996)

Stephen Guattery - Ph.D., Computer Science, Carnegie Mellon University, 1995. Computer Science [Parallel Numerical Algorithms, including Partitioning and Mapping]. (September 1995 to September 1997)

M. Ehtesham Hayder - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Fluid Mechanics [Computational Aeroacoustics]. (September 1995 to September 1997)

Victoria L. Interrante - Ph.D., Computer Science, University of North Carolina at Chapel Hill, 1996. Computer Science [Scientific Visualization]. (March 1996 to February 1998)

Leland M. Jameson - Ph.D., Applied Mathematics, Brown University, 1993. Applied & Numerical Mathematics [Multiresolution Schemes]. (August 1996 to March 1997)

Jim E. Jones - Ph.D., Applied Mathematics, University of Colorado-Boulder, 1995. Computer Science [Parallel Multigrid Methods]. (March 1995 to March 1997)

Kwan-Liu Ma - Ph.D., Computer Science, University of Utah, 1993. Computer Science [Visualization]. (May 1993 to August 1999)

Can Ozturan - Ph.D., Computer Science, Rensselaer Polytechnic Institute, 1995. Computer Science [System Software/Parallel Numerical Algorithms]. (August 1995 to August 1996)

J. Ray Ristorcelli - Ph.D., Mechanical and Aerospace Engineering, Cornell University, 1991. Fluid Mechanics [Turbulence Modeling]. (December 1992 to September 1996)

Moulay D. Tidriri - Ph.D., Fluid Mechanics, University of Paris IX, France, 1992. Computer Science [Domain Decomposition Methods and Parallel Solution of Partial Differential Equations]. (September 1994 to August 1996)

Linda F. Wilson - Ph.D., Electrical and Computer Engineering, University of Texas at Austin, 1994. Computer Science [Investigating the Issues of Object Migration in Object-Oriented Performance Tools and Object-Oriented Modeling of Architectures for Joint Performance/Reliability Analysis]. (August 1994 - July 1996)

VI. SENIOR SYSTEMS ANALYST

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to August 1997)

VII. VISITING SCIENTISTS

Mikhail M. Gilinsky - Ph.D., Aerodynamics, Moscow State University, 1965. Fluid Mechanics [Acoustics and Multidisciplinary Design Optimization]. (March 1996 to May 1996)

VIII. SHORT TERM VISITING SCIENTISTS

Saul Abarbanel - Ph.D., Theoretical Aerodynamics, Massachusetts Institute of Technology, 1959. Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics. (July to August 1996)

Ramesh K. Agarwal - Ph.D., Aeronautical Science, Stanford University, 1975. Bloomfield Distinguished Professor & Chair, Department of Aerospace Engineering, Wichita State University. Applied & Numerical Mathematics. (August 1996)

S. Balachandar - Ph.D., Engineering, Brown University, 1988. Assistant Professor, Department of Theoretical and Applied Mechanics, University of Illinois at Urbana. Applied & Numerical Mathematics. (July 1996)

Francoise Bataille - Ph.D., Compressible Turbulence, Ecole Centrale de Lyon, France, 1994. Associate Professor, Institut des Sciences Appliquees, Centre de Thermique, France. Fluid Mechanics. (July to August 1996)

Stanley A. Berger - Ph.D., Applied Mathematics, Brown University, 1959. Professor of Engineering Science, Department of Mechanical Engineering, University of California-Berkeley. Fluid Mechanics. (August 1996)

Christina L. Bloebaum - Ph.D., Aerospace Engineering, University of Florida, 1991. Assistant Professor, Department of Mechanical & Aerospace Engineering, State University of New York-Buffalo. Applied & Numerical Mathematics. (August 1996)

Shahid Bokhari - Ph.D., Electrical and Computer Engineering, University of Massachusetts-Amherst, 1978. Professor, Department of Electrical Engineering, University of Engineering & Technology, Lahore, Pakistan. Computer Science. (June to September 1996)

Andrew J. Booker - Ph.D., Mathematics, University of Washington, 1986. Statistician, Applied Statistics Department, The Boeing Company. Approximation Methods. (August 1996)

Michael Booty - Ph.D., Mathematics, Trinity College, University of Cambridge, 1978. Associate Professor, Department of Mathematics, New Jersey Institute of Technology. Applied & Numerical Mathematics. (May 1996)

Achi Brandt - Ph.D., Mathematics, Weizmann Institute of Science, 1965. Professor, Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Israel. Applied & Numerical Mathematics. (July and August 1996)

Denis R. Caromel - Ph.D., Computer Science, University of Nancy, France, 1991. Associate Professor, Department of Computer Science, University of Nice, France. Computer Science. (July to August 1996)

Nikos Chrisochoides - Ph.D., Computer Science, Purdue University, 1992. Research Assistant Professor, Advanced Computer Research Institute, Cornell Theory Center, Cornell University. Computer Science. (July to August 1996)

Timothy Clark - Ph.D., Mechanical Engineering, University of New Mexico, 1991. Staff Scientist, Theoretical Division, Los Alamos National Laboratory. Fluid Mechanics. (August 1996)

William O. Criminale - Ph.D., Aeronautics, The Johns Hopkins University, 1960. Professor, Department of Applied Mathematics, University of Washington. Fluid Mechanics. (June to July 1996)

Michel C. Delfour - Ph.D., Mathematics, Case Western Reserve University of Cleveland, 1970. Professor, Centre de Recherches Mathematiques, University of Montreal, Canada. Applied & Numerical Mathematics. (August 1996)

Julian A. Domaradzki - Ph.D., Physics, University of Warsaw, Poland, 1972. Associate Professor, Aerospace Engineering Department, University of Southern California. Fluid Mechanics. (June 1996)

Pierre R. Emeric - Ph.D., Applied Science, The College of William & Mary, 1995. Visiting Assistant Professor, Department of Mathematics, North Carolina State University. Applied & Numerical Mathematics [Methods for Determining Damage in Elastic Structures using Vibration Analysis]. (September 1996)

William W. Follett - M.S., Mechanical Engineering, Stanford University, 1987. Member of Technical Staff, CFD Technology Center, Design Technology Department, Rocketdyne Division of Rockwell International. Applied & Numerical Mathematics. (August 1996)

Daniele Funaro - Ph.D., Numerical Analysis, University of Pavia, Italy, 1983. Professor, Dipartimento Di Matematica, Universita Di Modena, Italy. Applied & Numerical Mathematics. (July to August 1996)

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Computational Methods for Partial Differential Equations]. (July and August 1996)

Mohamed M. Hafez - Ph.D., Aeronautics, University of Southern California, 1972. Professor, Department of Mechanical and Aeronautical Engineering, University of California-Davis. Applied & Numerical Mathematics. (August 1996)

Maurice Holt - Ph.D., Mathematics, University of Manchester, England, 1948. Professor, College of Engineering, Mechanical Engineering, University of California-Berkeley. Applied & Numerical Mathematics [Godunov Methods]. (July 1996)

Graham A. Horton - Ph.D., Computer Science, Universitat Erlangen-Nurnberg, Germany, 1991. Research/Lecturer, Universitat Erlangen-Nurnberg, Germany. Computer Science. (August 1996)

Donald R. Jones - Ph.D., Economics, Massachusetts Institute of Technology, 1984. Staff Research Scientist, Consumer and Operations Research Department, General Motors Research and Development Center. Approximation Methods. (August 1996)

Harry F. Jordan - Ph.D., Physics, University of Illinois, 1967. Professor, Department of Electrical and Computer Engineering, University of Colorado-Boulder. Computer Science [Parallel Processing Systems]. (June 1996)

Carl Kesselman - Ph.D., Computer Science, University of California-Los Angeles, 1991. Member of The Beckman Institute, California Institute of Technology. Computer Science. (July to August 1996)

James R. Koehler - Ph.D., Statistics, Stanford University, 1990. Assistant Professor, Mathematics Department, University of Colorado-Denver. Approximation Methods. (August 1996)

Koen G. Langendoen - Ph.D., Computer Science, Universiteit van Amsterdam, 1993. Assistant Professor, Department of Mathematics and Computer Science, Vrije Universiteit, Amsterdam, Holland. Computer Science. (May 1996)

Rainald Lohner - Ph.D., Civil Engineering, University of Swansea, United Kingdom, 1985. Professor, Institute for Computational Sciences and Informatics, George Mason University. Computer Science. (August 1996)

Alex Mahalov - Ph.D., Applied Mathematics, Cornell University, 1991. Assistant Professor, Department of Mathematics, Arizona State University. Fluid Mechanics. (May to June 1996)

Stephen F. McCormick - Ph.D., Mathematics, University of Southern California, 1971. Professor, Program of Applied Mathematics, University of Colorado. Applied & Numerical Mathematics. (May 1996)

Amar Mukherjee - Ph.D., Physics, University of Calcutta, 1962. Professor, Department of Computer Science, University of Central Florida. Computer Science. (July 1996)

Anthony T. Patera - Ph.D., Applied Mathematics, Massachusetts Institute of Technology, 1982. Professor, Department of Mechanical Engineering, Massachusetts Institute of Technology. Applied & Numerical Mathematics. (August 1996)

Luca Pavarino - Ph.D., Mathematics, Courant Institute, New York University, 1992. Assistant Professor, Department of Mathematics, University of Pavia; Italy. Applied & Numerical Mathematics. (August to September 1996)

Yousef Saad - Ph.D., Computer Science, University of Grenoble, France, 1983. Professor, Department of Computer Science, University of Minnesota. Computer Science. (August 1996)

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics. (May and July 1996)

Arun K. Somani - Ph.D., Computer Engineering, McGill University, 1985. Professor, Department of Electrical Engineering, University of Washington. Computer Science. (August 1996)

Charles G. Speziale - Ph.D., Aerospace and Mechanical Sciences, Princeton University, 1978. Professor, Aerospace & Mechanical Engineering Department, Boston University. Fluid Mechanics. (August 1996)

S.S. Sritharan - Ph.D., Applied Mathematics, University of Arizona, 1982. Interim University Professor, Department of Science and Technology, Naval Command Control & Ocean Surveillance Center-San Diego. Fluid Mechanics. (June to July 1996)

Xian-He Sun - Ph.D., Computer Science, Michigan State University, 1990. Assistant Professor, Department of Computer Science, Louisiana State University. Computer Science. (July to August 1996)

Siva Thangam - Ph.D., Mechanical Engineering, Rutgers University, 1980. Professor, Department of Mechanical Engineering, Stevens Institute of Technology. Fluid Mechanics. (June to July 1996)

Michael W. Trosset - Ph.D., Statistics, University of California-Berkeley, 1983. Adjunct Lecturer, Department of Mathematics, University of Arizona. Approximation Methods. (August 1996)

Eli Turkel - Ph.D., Applied Mathematics, New York University, 1970. Associate Professor, Department of Applied Mathematics, Tel-Aviv University, Israel. Applied & Numerical Mathematics. (July to August 1996)

Layne Watson - Ph.D., Mathematics, University of Michigan, 1974. Professor, Departments of Computer Science and Mathematics, Virginia Polytechnic Institute and State University. Approximation Methods. (August 1996)

Olof B. Widlund - Ph.D., Computer Science, Uppsala University, Sweden, 1966. Professor, Department of Computer Science and Mathematics, New York University. Computer Science. (August 1996)

Stephen L. Woodruff - Ph.D., Aerospace Engineering, The University of Michigan, 1986. Visiting Researcher, Center for Fluid Mechanics, Brown University. Fluid Mechanics. (May to June 1996)

Kun Xu - Ph.D., Astrophysics, Columbia University, 1993. Research Staff, Department of Mechanical and Aerospace Engineering, Princeton University. Applied & Numerical Mathematics. (June 1996)

P.K. Yeung - Ph.D., Mechanical Engineering, Cornell University, 1989. Assistant Professor, School of Aerospace Engineering, Georgia Institute of Technology. Fluid Mechanics. (June 1996)

Bassam Younis - Ph.D., Mechanical Engineering-Fluid Mechanics, Imperial College, London, England, 1984. Senior Lecturer, Department of Civil Engineering, City University, London, England. Fluid Mechanics [Turbulence Modeling]. (June to August 1996)

Hong Zhang-Sun - Ph.D., Applied Mathematics, Michigan State University, 1989. Visiting Scientist, Department of Mathematical Sciences, Clemson University. Applied & Numerical Mathematics. (July to August 1996)

Mei Zhuang - Ph.D., Aeronautics, California Institute of Technology, 1990. Assistant Professor, Department of Mechanical Engineering, Michigan State University. Fluid Mechanics. (August 1996)

IX. SENIOR RESEARCH ASSOCIATE

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

X. CONSULTANTS

Ivo Babuska - Ph.D., Technical Science, Technical University, Prague, Czechoslovakia, 1951; Mathematics, Academy of Science, Prague, 1956; D.Sc., Mathematics, Academy of Science, Prague, 1960. Robert Trull Chair in Engineering, TICAM, The University of Texas at Austin. Applied & Numerical Mathematics [Finite Element Methods Associated With Structural Engineering]

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

David Banks - Ph.D., Computer Science, University of North Carolina, 1993. Assistant Professor, Department of Computer Science, Mississippi State University. Computer Science [Graphics and Visualization]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Alvin Bayliss - Ph.D., Mathematics, New York University, 1975. Associate Professor, Technological Institute, Northwestern University. Fluid Mechanics [Numerical Solution of the Equations of Fluid Flow and Acoustics]

John A. Burns - Ph.D., Mathematics, University of Oklahoma, 1973. Professor, Virginia Polytechnic Institute and State University. Applied & Numerical Mathematics [Numerical Methods in Optimal Design and Control]

Xiao-Chuan Cai - Ph.D., Mathematics, Courant Institute of Mathematical Science, New York University, 1989. Assistant Professor, Department of Computer Science, University of Colorado-Boulder. Computer Science [Numerical Analysis and Parallel Computing]

Tony F. Chan - Ph.D., Computer Science, Stanford University, 1978. Professor, Department of Mathematics, University of California-Los Angeles. Applied & Numerical Mathematics [Multigrid and Domain Decomposition for Unstructured Grid]

Barbara M. Chapman - M.S., Mathematics, University of Canterbury, Christchurch, New Zealand, 1985. Director, European Institute for Parallel Computing, University of Vienna. Computer Science [Parallel Language Extensions and Optimizations for Parallel Compilers]

Gianfranco Ciardo - Ph.D., Computer Science, Duke University, 1989. Assistant Professor, The College of William & Mary. Computer Science [Reliability Models]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Geoffrey Fox - Ph.D., Physics, Cambridge University, 1967. Professor, Department of Computer Science, Syracuse University. Computer Science [Networking]

Dennis B. Gannon - Ph.D., Mathematics, University of California, Davis, 1974. Professor, Department of Computer Science, Indiana University-Bloomington. Computer Science [Investigation of Algorithms and Programming Techniques for Parallel Computers]

James F. Geer - Ph.D., Applied Mathematics, New York University, 1967. Professor, Systems Science and Mathematical Sciences, Watson School of Engineering, Applied Science and Technology, SUNY-Binghamton. Applied & Numerical Mathematics [Approximation Method of Solutions for Partial Differential Equations]

Mikhail M. Gilinsky - Ph.D., Aerodynamics, Moscow State University, 1965. Research Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Theoretical and Numerical Research on Nozzle-Jet Flows including Turbulence Models, Acoustic Field Calculation, and Numerical Simulation of Nozzle Flows]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Professor, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Max Gunzburger - Ph.D., Applied Mathematics, New York University, 1969. Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Numerical Methods for Flow Control Problems]

Matthew D. Haines - Ph.D., Computer Science, Carnegie Mellon University, 1995. Assistant Professor, Department of Computer Science, University of Wyoming. Computer Science [Parallel Programming Environment and Run Time Systems]

Gene J.-W. Hou - Ph.D., Computational Mechanics, Design Optimization, University of Iowa, 1983. Associate Professor, Mechanical Engineering Department, Old Dominion University. Applied & Numerical Mathematics [Computational Mechanics Design Optimization]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Instability and Transition]

Hideaki Kaneko - Ph.D., Applied Functional Analysis, Clemson University, 1979. Professor, Department of Mathematics and Statistics, Old Dominion University. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]

Ashwani K. Kapila - Ph.D., Theoretical and Applied Mechanics, Cornell University, 1975. Associate Professor, Department of Mathematical Sciences, Rensselaer Polytechnic Institute. Fluid Mechanics [Mathematical Combustion]

Ken Kennedy - Ph.D., Computer Science, New York University, 1971. Chairman, Department of Computer Science, Rice University. Computer Science [Parallel Compilers and Languages]

Charles Koelbel - Ph.D., Computer Science, Purdue University, 1990. Research Scientist, Department of Computer Science, Rice University. Computer Science [Compilers for Parallel Computers]

David A. Kopriva - Ph.D., Applied Mathematics, University of Arizona, 1982. Associate Professor, Department of Mathematics and SCRI, Florida State University. Applied & Numerical Mathematics [Development and Implementation of Spectral Methods for Viscous Compressible Flows]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics

Heinz-Otto Kreiss - Ph.D., Mathematics, Royal Institute of Technology, Sweden, 1960. Professor, Department of Applied Mathematics, California Institute of Technology, Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

Vipin Kumar - Ph.D., Computer Science, University of Maryland, 1982. Associate Professor, Department of Computer Science, University of Minnesota. Computer Science [Parallel Sparse Matrix Algorithms and Applications]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

Scott T. Leutenegger - Ph.D., Computer Science, University of Wisconsin-Madison, 1990. Assistant Professor, Department of Mathematics and Computer Science, University of Denver. Computer Science [System Software Related to Databases for Scientific Data]

Robert W. MacCormack - M.S., Mathematics, Stanford University. Professor, Department of Aeronautics and Astronautics, Stanford University. Fluid Mechanics [Numerical Methods for the Solution of the Equations of Fluid Mechanics]

Kurt Maly - Ph.D., Computer Science, Courant Institute, New York University, 1973. Kaufman Professor and Chair, Department of Computer Science, Old Dominion University. Computer Science [High Performance Communication]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Sanjoy K. Mitter - Ph.D., Electrical Engineering, Imperial College of Science & Technology, London, 1965. Professor of Electrical Engineering, Co-Director, Laboratory for Information and Decision Systems, Director, Center for Intelligent Control Systems, Massachusetts Institute of Technology. Fluid Mechanics [Control Theory]

David M. Nicol - Ph.D., Computer Science, University of Virginia, 1985. Professor, Department of Computer Science, College of William & Mary. Computer Science [Mapping Algorithms onto Parallel Computing Systems]

R.A. Nicolaides - Ph.D., Computer Science, University of London, 1972. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Solution of Partial Differential Equations]

Demetrius Papageorgiou - Ph.D., Mathematics, University of London, 1985. Assistant Professor, Department of Mathematics, New Jersey Institute of Technology. Fluid Mechanics [Theoretical and Computational Fluid Dynamics]

Alex Pothén - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Daniel A. Reed - Ph.D., Computer Science, Purdue University, 1983. Assistant Professor, Department of Computer Science, University of Illinois. Computer Science [Performance Evaluation of Computer Systems]

John H. Reif - Ph.D., Applied Mathematics, Harvard University, 1977. Professor, Department of Computer Science, Duke University. Computer Science [Parallel Algorithms]

Philip L. Roe - Ph.D., Aeronautics, University of Cambridge, United Kingdom, 1962. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Numerical Mathematics and Aeroacoustics]

Ahmed H. Sameh - Ph.D., Civil Engineering, University of Illinois, 1968. Head, William Norris Chair, and Professor, Department of Computer Science, University of Minnesota. Computer Science [Numerical Algorithms]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Charles G. Speziale - Ph.D., Aerospace and Mechanical Sciences, Princeton University, 1978. Professor, Aerospace & Mechanical Engineering Department, Boston University. Fluid Mechanics [Turbulence Modeling]

Shlomo Ta'asan - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1985. Professor, Department of Mathematics, Carnegie Mellon University. Applied & Numerical Mathematics [Numerical Analysis and Algorithm Development]

Lu Ting - Ph.D., Aeronautics, New York University, 1951. Professor, Courant Institute of Mathematical Sciences, New York University. Fluid Mechanics [Nonlinear Acoustic/Structure Interaction Problems]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization including Multidisciplinary Design Optimization]

Kishor Trivedi - Ph.D., Computer Science, University of Illinois-Urbana, 1974. Professor, Department of Electrical Engineering, Duke University. Computer Science [Performance and Reliability Modeling Methods, Tools and Applications]

George M. Vahala - Ph.D., Physics, University of Iowa, 1972. Professor, Department of Physics, The College of William & Mary. Fluid Mechanics [Group Renormalization Methods for Turbulence Approximation]

Bram van Leer - Ph.D., Theoretical Astrophysics, Leiden State University, The Netherlands, 1970. Professor, Department of Aerospace Engineering, University of Michigan. Applied & Numerical Mathematics [Computational Fluid Dynamics]

Linda F. Wilson - Ph.D., Electrical and Computer Engineering, University of Texas-Austin, 1994. Professor, Thayer School of Engineering, Dartmouth College. Computer Science [Parallel Discrete-Event Simulation]

Hans Zima - Ph.D., Mathematics, University of Vienna, Austria, 1964. Professor, Institute for Software Technology and Parallel Systems, University of Vienna, Austria. Computer Science [Compiler Development for Parallel and Distributed Multiprocessors]

XI. GRADUATE STUDENTS

Abdelkader Baggag - Department of Computer Science, The University of Minnesota. (September 1995 to Present)

Bendiks J. Boersma - Lab for Aerio- & Hydrodynamics, Delft University of Technology. (July to August 1996)

David C. Cronk - Department of Computer Science, The College of William & Mary. (August 1993 to Present)

Indraneel Das - Department of Computational & Applied Mathematics, Rice University. (May to July 1996)

Adi Ditkowski - School of Mathematical Sciences, Tel-Aviv University. (July to August 1996)

Dawn M. Galayda - Department of Computer Science, The College of William & Mary. (September 1995 to April 1996)

Jeffrey A. Hittinger - Department of Aerospace Engineering, The University of Michigan. (May to August 1996)

Nilan Karunaratne - Department of Computer Science, Old Dominion University. (August 1995 to Present)

Joe L. Manthey - Department of Mathematics, Old Dominion University. (September 1993 to Present)

Deborah F. Pilkey - Department of Engineering Sciences & Mechanics, Virginia Polytechnic Institute and State University. (October 1995 to Present)

Kevin Roe - Department of Computer Science, The College of William & Mary. (May 1995 to Present)

Robert V. Wilson - Department of Mechanical Engineering and Mechanics, Old Dominion University. (October 1992 to Present)

Chuan-Kai Yang - Computer Science Department, State University of New York at Stony Brook. (June to August 1996)

XII. STUDENT ASSISTANTS

Kathryn Paulson - Student at Ferguson High School. (June to August 1996)

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE November 1996	3. REPORT TYPE AND DATES COVERED Contractor Report		
4. TITLE AND SUBTITLE Semiannual Report. April 1, 1996 through September 30, 1996		5. FUNDING NUMBERS C NAS1-19480 WU 505-90-52-01		
6. AUTHOR(S)				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Computer Applications in Science and Engineering Mail Stop 403, NASA Langley Research Center Hampton, VA 23681-0001		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Langley Research Center Hampton, VA 23681-0001		10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA CR-201622		
11. SUPPLEMENTARY NOTES Langley Technical Monitor: Dennis M. Bushnell Final Report				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category 59		12b. DISTRIBUTION CODE		
13. ABSTRACT (Maximum 200 words) This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, fluid mechanics, and computer science during the period April 1, 1996 through September 30, 1996.				
14. SUBJECT TERMS Applied Mathematics; Numerical Analysis; Fluid Mechanics; Computer Science			15. NUMBER OF PAGES 90	
			16. PRICE CODE A05	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	